

(12) UK Patent Application (19) GB (11) 2 349 252 (13) A

(43) Date of A Publication 25.10.2000

(21) Application No 0018379.8

(22) Date of Filing 10.06.1997

Date Lodged 26.07.2000

(30) Priority Data

(31) 08661027 (32) 10.06.1996 (33) US

(31) 08661003 (32) 10.06.1996

(31) 08659709 (32) 10.06.1996

(31) 08659708 (32) 10.06.1996

(62) Divided from Application No 9825726.4 under Section 15(4) of the Patents Act 1977

(71) Applicant(s)

LSI Logic Corporation

(Incorporated in USA - Delaware)

1551 McCarthy Blvd, Milpitas, California 95035,
United States of America

(72) Inventor(s)

Frank Worrell

Hartwig Ekner

(51) INT CL⁷

G06F 9/30 9/318

(52) UK CL (Edition R)

G4A AJR APX

(56) Documents Cited

GB 2289353 A GB 2284492 A JP 070121352 A

US 4437149 A

14 Annual Int. Symp. on Computer Architecture, Conf.
Proc., 1987, IEEE, pp 309-319

(58) Field of Search

UK CL (Edition R) G4A AJR APX

INT CL⁷ G06F 9/30 9/318

(74) Agent and/or Address for Service

R G C Jenkins & Co

26 Caxton Street, LONDON, SW1H ORJ,
United Kingdom

(54) Abstract Title

An apparatus and method for detecting and decompressing instructions from a variable length compressed instruction set

(57) A microprocessor is configured to fetch a compressed instruction set which comprises a subset of a corresponding non-compressed instruction set. The compressed instruction set is a variable length instruction set including, for example 16-bit and 32-bit instructions. Extended (e.g. 32-bit) instructions are coded using an extend opcode, which indicates that the instruction being fetched is an extended (e.g. 32 bit) instruction. Extended instructions have a second opcode field, which is decompressed into a decompressed opcode field within a decompressed instruction. For non-extended instructions, the first opcode field is decompressed into the decompressed opcode field. The compressed instruction set further includes multiple sets of register mappings from the compressed register fields to the decompressed register fields. Certain select instructions are assigned two opcode encodings, one for each of two mappings of the corresponding register fields. The compressed register field is directly copied into a portion of the decompressed register field while the remaining portion of the decompressed register field is created using a small number of logic gates. The subroutine call instruction within the compressed instruction set includes a compression mode which indicates whether or not the target routine is coded in compressed instructions. The compression mode is stored in the program counter register. The decompression of the immediate field used for load/store instructions having the global pointer register as a base register is optimized for mixed compressed/non-compressed instruction execution. The immediate field is decompressed into a decompressed immediate field for which the most significant bit is set.

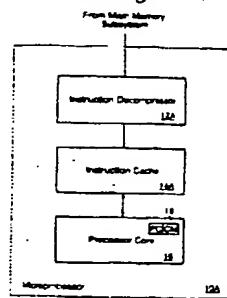


FIG. 1

GB 2 349 252 A

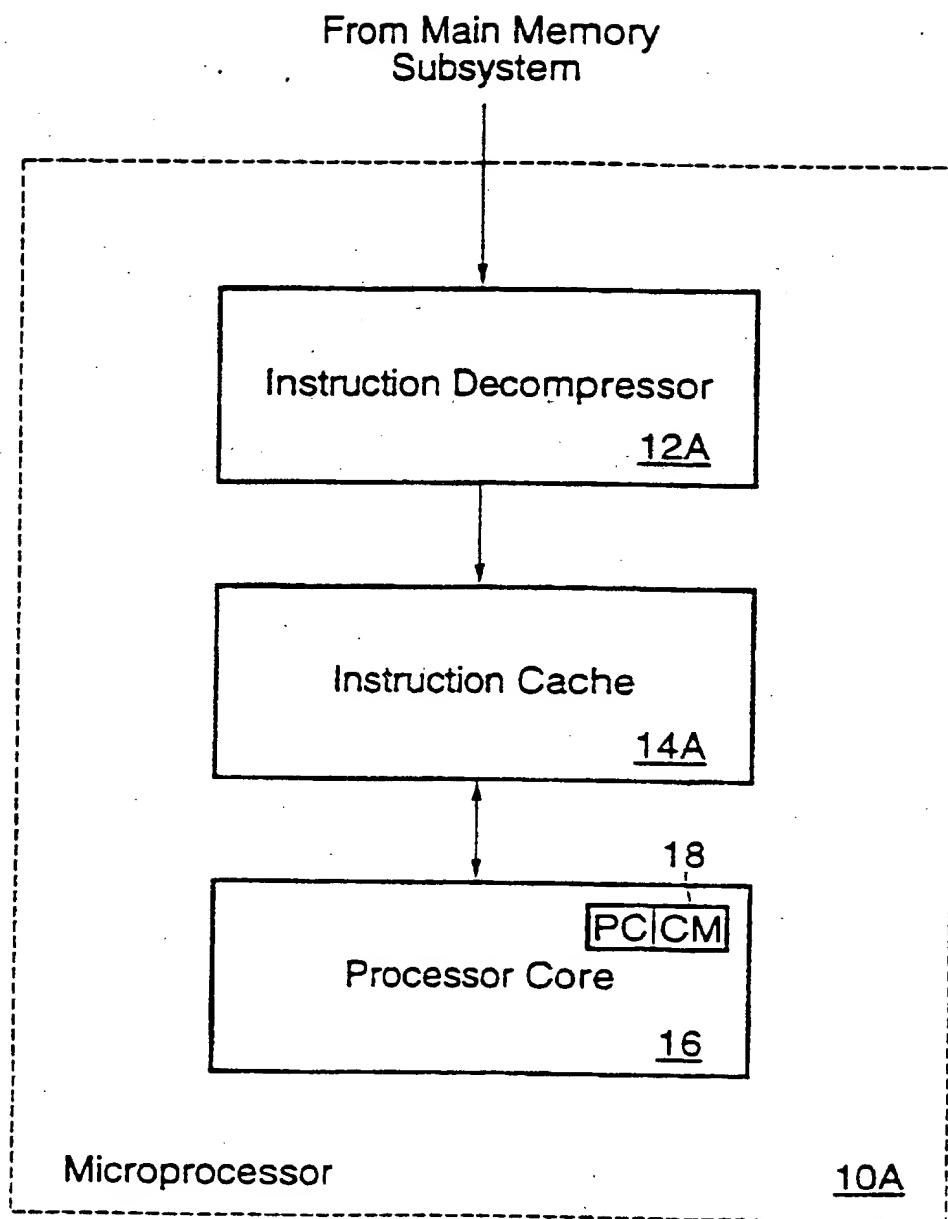


FIG. 1

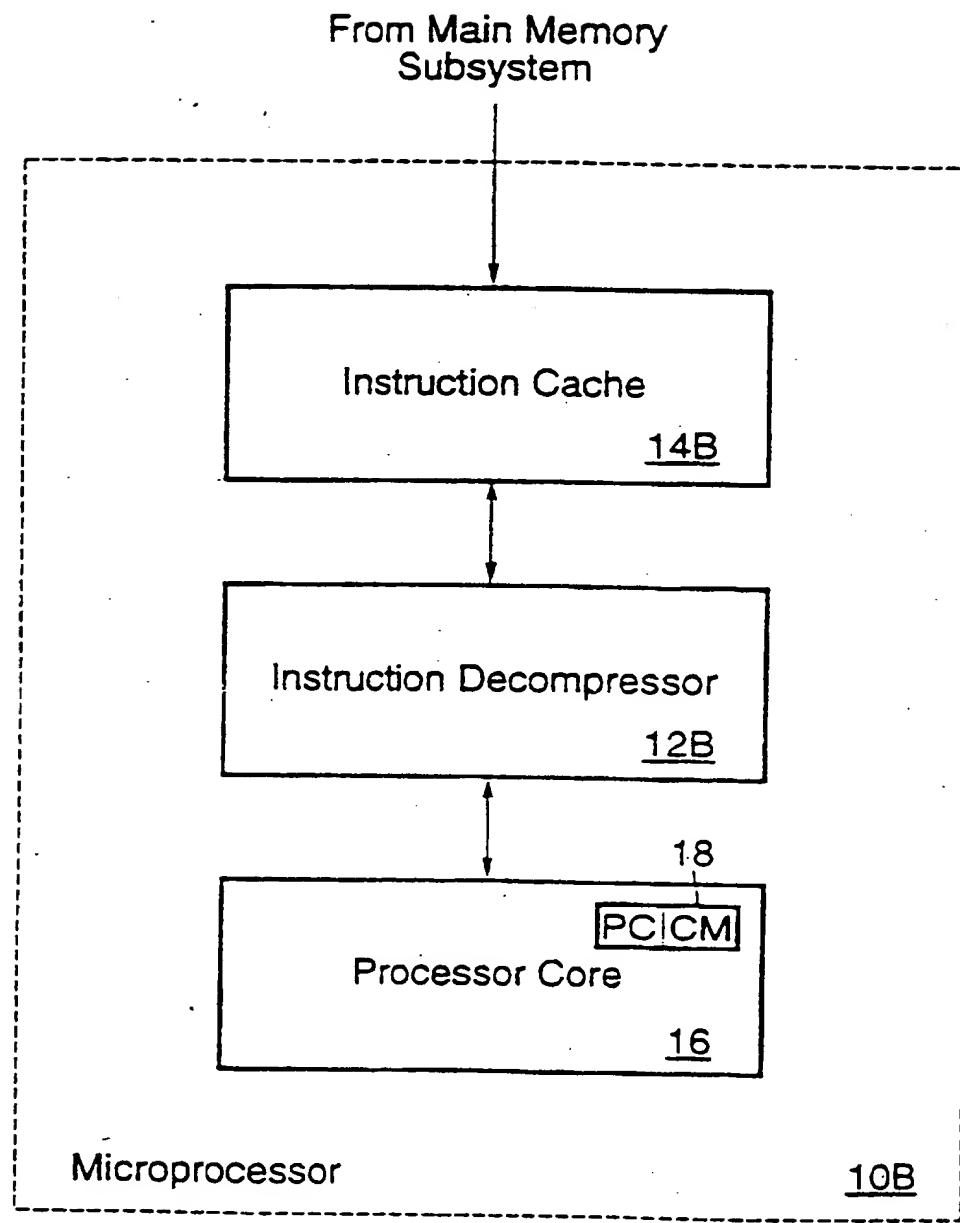


FIG. 2

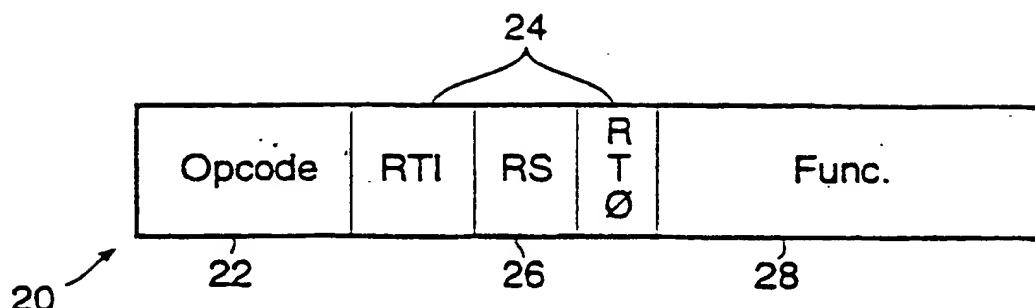


FIG. 3A

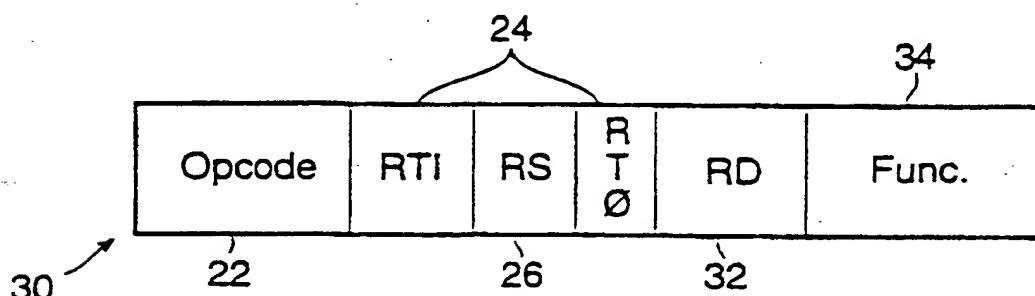


FIG. 3B

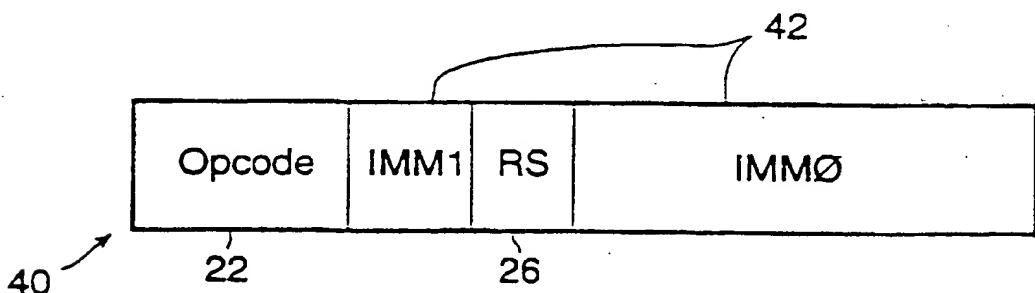


FIG. 3C

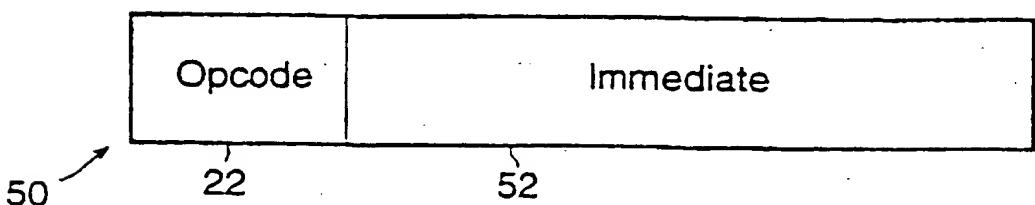


FIG. 3D

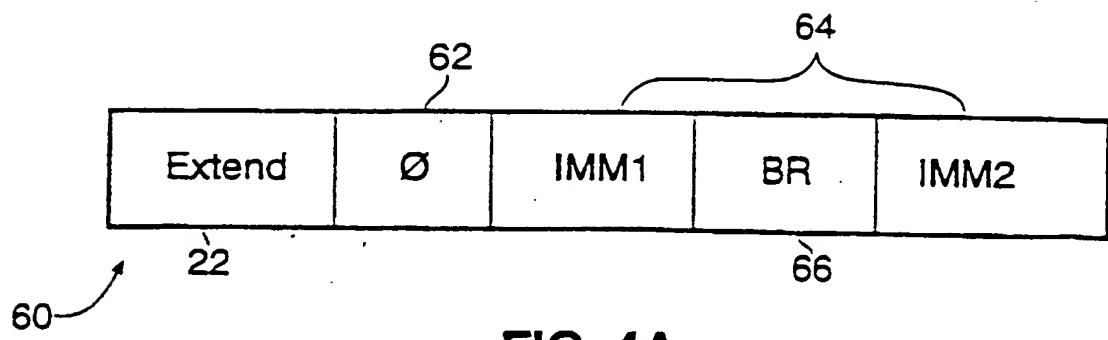


FIG. 4A

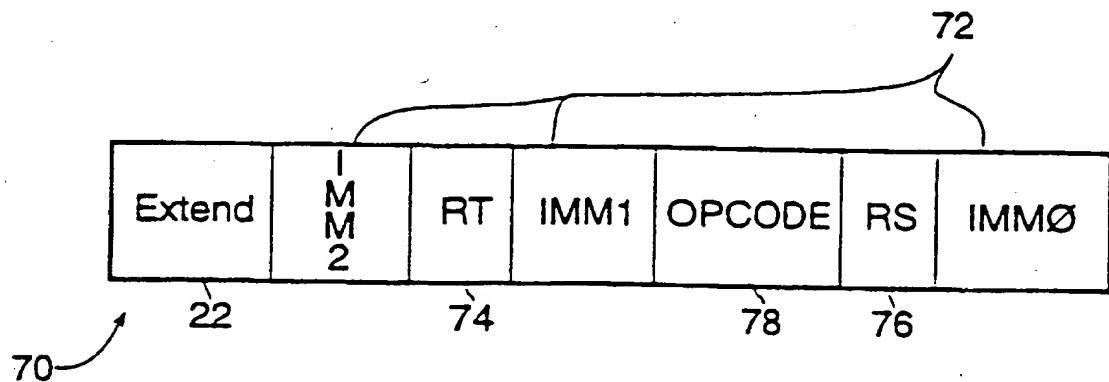


FIG. 4B

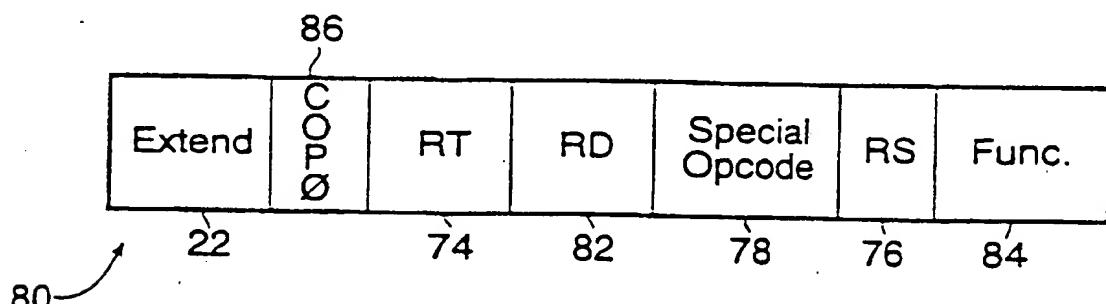


FIG. 4C

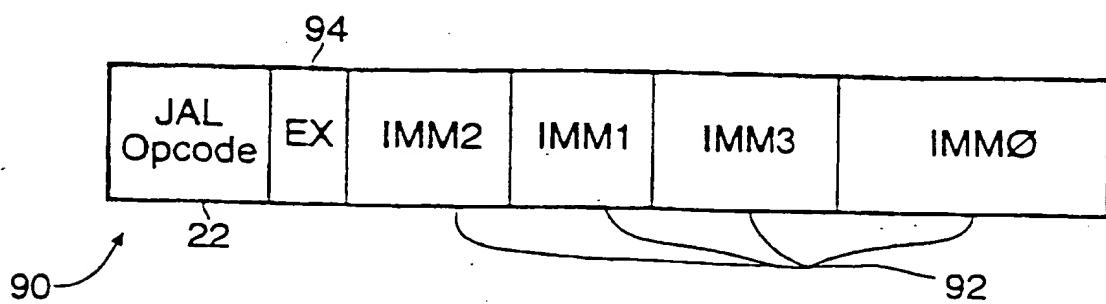


FIG. 4D

106

104

100

102

| <u>OPCODE</u> | <u>FUNC.</u> | <u>INSTRUCTION</u> | <u>OPERANDS</u> |
|---------------|--------------|--------------------|-----------------|
| 11 | imm5 | LH | rt, rs |
| 13 | imm5 | LW | rt, xs |
| 15 | imm5 | LHU | rt, rs |
| 17 | imm5 | LW | rt, rs |
| 18 | imm5 | LB | rt, rs |
| 19 | imm5 | SH | rt, rs |
| 1a | imm5 | SB | rt, rs |
| 1b | imm5 | SW | rt, xs |
| 1c | imm5 | LBU | rt, rs |
| 1d | | special | |
| 1f | imm5 | SW | rt, rs |
| 12 | 02 | SLT | rs, rt |
| 12 | 06 | SLTU | rs, rt |
| 12 | 08 | MOVE | rs, xt |
| 12 | 0a | SLT | rs, xt |
| 12 | 0c | ADDU | rs, xt |
| 12 | 0e | SLTU | rs, xt |
| 12 | 10 | MOVE | xs, rt |
| 12 | 12 | SLT | xs, rt |
| 12 | 14 | ADDU | xs, rt |
| 12 | 16 | SLTU | xs, rt |
| 12 | 18 | MOVE | xs, xt |
| 12 | 1a | SLT | xs, xt |
| 12 | 1c | ADDU | xs, xt |
| 12 | 1e | SLTU | xs, xt |

| 106 | 104 | 100 | 102 |
|---------------|--------------|--------------------|-----------------|
| <u>OPCODE</u> | <u>FUNC.</u> | <u>INSTRUCTION</u> | <u>OPERANDS</u> |
| 10 | 02 | SRL | rt, imm3 |
| 10 | 03 | SRA | rt, imm3 |
| 10 | 04 | SLLV | rt, rs |
| 10 | 06 | SRLV | rt, rs |
| 10 | 07 | SRAV | rt, rs |
| 10 | 08 | JR | rs |
| 10 | 09 | JALR | xt, rs |
| 10 | 0c | SYSCALL | imm6 |
| 10 | 0d | BREAK | imm6 |
| 10 | 10 | ADD | rs, rt |
| 10 | 12 | SUB | rs, rt |
| 10 | 13 | NEG | rs, rt |
| 10 | 14 | AND | rs, rt |
| 10 | 15 | OR | rs, rt |
| 10 | 16 | XOR | rs, rt |
| 10 | 17 | NOT | rs, rt |

112

FIG. 5B

| | 106 | 104 | 100 | 102 |
|---------------|--------------|--------------------|-----------------|-----|
| <u>OPCODE</u> | <u>FUNC.</u> | <u>INSTRUCTION</u> | <u>OPERANDS</u> | |
| 12 | 01 | ADDU | r0, rs, rt | |
| 12 | 03 | SUBU | r0, rs, rt | |
| 12 | 05 | ADDU | r1, rs, rt | |
| 12 | 07 | SUBU | r1, rs, rt | |
| 12 | 09 | ADDU | r2, rs, rt | |
| 12 | 0b | SUBU | r2, rs, rt | |
| 12 | 0d | ADDU | r3, rs, rt | |
| 12 | 0f | SUBU | r3, rs, rt | |
| 12 | 11 | ADDU | r4, rs, rt | |
| 12 | 13 | SUBU | r4, rs, rt | |
| 12 | 15 | ADDU | r5, rs, rt | |
| 12 | 17 | SUBU | r5, rs, rt | |
| 12 | 19 | ADDU | r6, rs, rt | |
| 12 | 1b | SUBU | r6, rs, rt | |
| 12 | 1d | ADDU | r7, rs, rt | |
| 12 | 1f | SUBU | r7, rs, rt | |
| 16 | imm | SLL | rd, rt, imm | |

FIG. 5C

| <u>OPCODE</u> | <u>INSTRUCTION</u> | <u>OPERANDS</u> |
|---------------|--------------------|-----------------|
| 00 | BEQZ | xs, simm8 |
| 01 | BNEZ | xs, simm8 |
| 04 | BEQZ | rs, simm8 |
| 05 | BNEZ | rs, simm8 |
| 08 | MOVEI | rs, imm8 |
| 09 | ADDIU | rs, simm8 |
| 0a | SLTI | rs, imm8 |
| 0b | SLTIU | rs, imm8 |
| 0c | ANDI | rs, imm8 |
| 0e | CMPI | rs, imm8 |

116

FIG. 5D

| <u>OPCODE</u> | <u>INSTRUCTION</u> | <u>OPERANDS</u> |
|---------------|--------------------|-----------------|
| 02 | BR | simm11 |
| 0f | extend | imm11 |

118

FIG. 5E

| <u>OPCODE</u> | <u>INSTRUCTION</u> | <u>OPERANDS</u> |
|---------------|--------------------|-----------------|
| 04 | BEQ | rs, rt, simm12 |
| 05 | BNE | rs, rt, simm12 |
| 06 | BLEZ | rs, simm12 |
| 07 | BGTZ | rs, simm12 |
| 08 | ADDI | rt, rs, simm12 |
| 09 | ADDIU | rt, rs, simm12 |
| 0a | SLTI | rt, rs, simm12 |
| 0b | SLTIU | rt, rs, simm12 |
| 0c | ANDI | rt, rs, simm12 |
| 0d | ORI | rt, rs, simm12 |
| 0e | XORI | rt, rs, simm12 |
| 0f | LUI | rt, simm12 |
| 10 | LB | rt, rs, simm12 |
| 11 | LH | rt, rs, simm12 |
| 12 | LWL | rt, rs, simm12 |
| 13 | LW | rt, rs, simm12 |
| 14 | LBU | rt, rs, simm12 |
| 15 | LHU | rt, rs, simm12 |
| 16 | LWR | rt, rs, simm12 |
| 18 | SB | rt, rs, simm12 |
| 19 | SH | rt, rs, simm12 |
| 1a | SWL | rt, rs, simm12 |
| 1b | SW | rt, rs, simm12 |
| 1e | SWR | rt, rs, simm12 |

| <u>RT</u> | <u>INSTRUCTION</u> | <u>OPERANDS</u> |
|-----------|--------------------|-----------------|
| 00 | BLTZ | rs, simm12 |
| 01 | BGEZ | rs, simm12 |
| 02 | BLTZL | rs, simm12 |
| 03 | BGEZL | rs, simm12 |
| 08 | TGEI | rs, simm12 |
| 09 | TGEIU | rs, simm12 |
| 0a | TLTI | rs, simm12 |
| 0b | TLTIU | rs, simm12 |
| 0c | TEQI | rs, simm12 |
| 0e | TNEI | rs, simm12 |
| 10 | BLTZAL | rs, simm12 |
| 11 | BGEZAL | rs, simm12 |
| 12 | BLTZALL | rs, simm12 |
| 13 | BGEZALL | rs, simm12 |

122

FIG. 6B

| <u>FUNC.</u> | <u>INSTRUCTION</u> | <u>OPERANDS</u> |
|--------------|--------------------|-----------------|
| 00 | SLL | rd, rt, imm5 |
| 02 | SRL | rd, rt, imm5 |
| 03 | SRA | rd, rt, imm5 |
| 04 | SLLV | rd, rt, rs |
| 06 | SRLV | rd, rt, rs |
| 07 | SRAV | rd, rt, rs |
| 08 | JR | rs |
| 09 | JALR | rd, rs |
| 0c | SYSCALL | imm15 |
| 0d | BREAK | imm15 |
| 0f | SYNC | |
| 10 | MFHI | rd |
| 11 | MTHI | rs |
| 12 | MFLO | rd |
| 13 | MTLO | rs |
| 18 | MULT | rs, rt |
| 19 | MULTU | rs, rt |
| 1a | DIV | rs, rt |
| 1b | DIVU | rs, rt |

124

FIG. 6C

12 / 20

| <u>FUNC.</u> | <u>INSTRUCTION</u> | <u>OPERANDS</u> |
|--------------|--------------------|-----------------|
| 20 | ADD | rd, rs, rt |
| 21 | ADDU | rd, rs, rt |
| 22 | SUB | rd, rs, rt |
| 23 | SUBU | rd, rs, rt |
| 24 | AND | rd, rs, rt |
| 25 | OR | rd, rs, rt |
| 26 | XOR | rd, rs, rt |
| 27 | NOR | rd, rs, rt |
| 2a | SLT | rd, rs, rt |
| 2b | SLTU | rd, rs, rt |
| 30 | TGE | rs, rt |
| 31 | TGEU | rs, rt |
| 32 | TLT | rs, rt |
| 33 | TLTU | rs, rt |
| 34 | TEQ | rs, rt |
| 36 | TNE | rs, rt |

126

FIG. 6D

| <u>RS, RT.</u> | <u>INSTRUCTION</u> | <u>OPERANDS</u> |
|----------------|--------------------|-----------------|
| 00,* | MFC0 | rt, rd |
| 02,* | CFC0 | rt, rd |
| 04,* | MTC0 | rt, rd |
| 06,* | CTC0 | rt, rd |
| 08, 00 | BC0F | imm6 |
| 08, 01 | BC0T | imm6 |
| 08, 02 | BC0FL | imm6 |
| 08, 03 | BC0TL | imm6 |
| 1x,* | COP0 | Func. (Fig. 6F) |

126



FIG. 6E

| <u>FUNC.</u> | <u>INSTRUCTION</u> |
|--------------|--------------------|
| 01 | TLBR |
| 02 | TLBWI |
| 06 | TLBWR |
| 08 | TLBP |
| 10 | RFE |
| 18 | ERET |

130



FIG. 6F

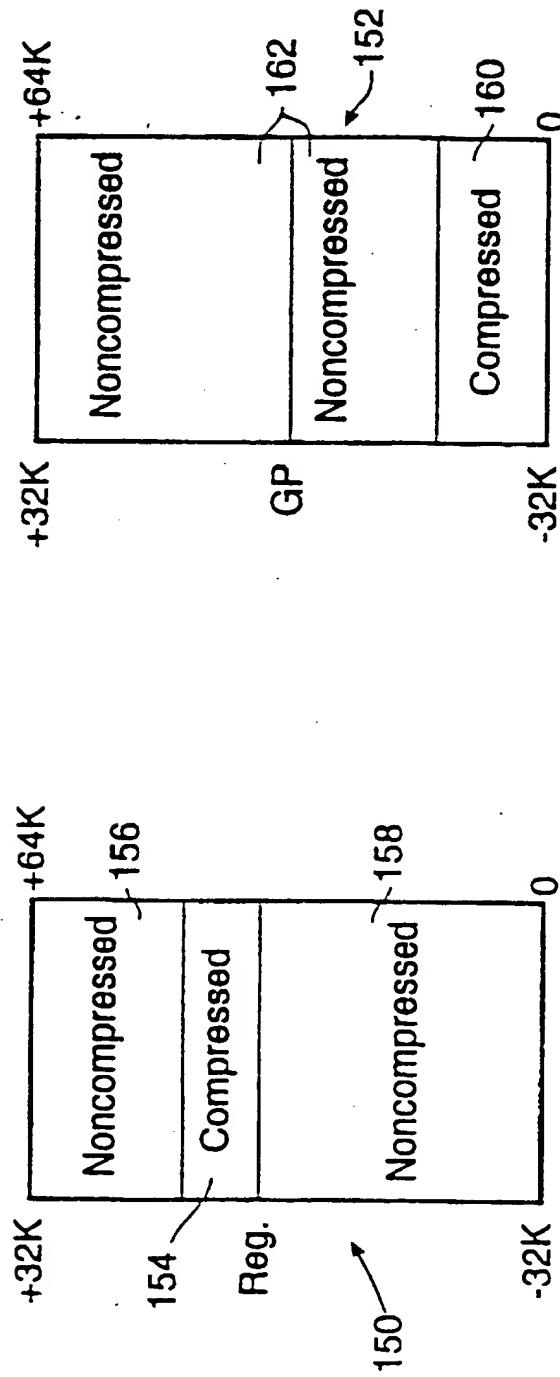
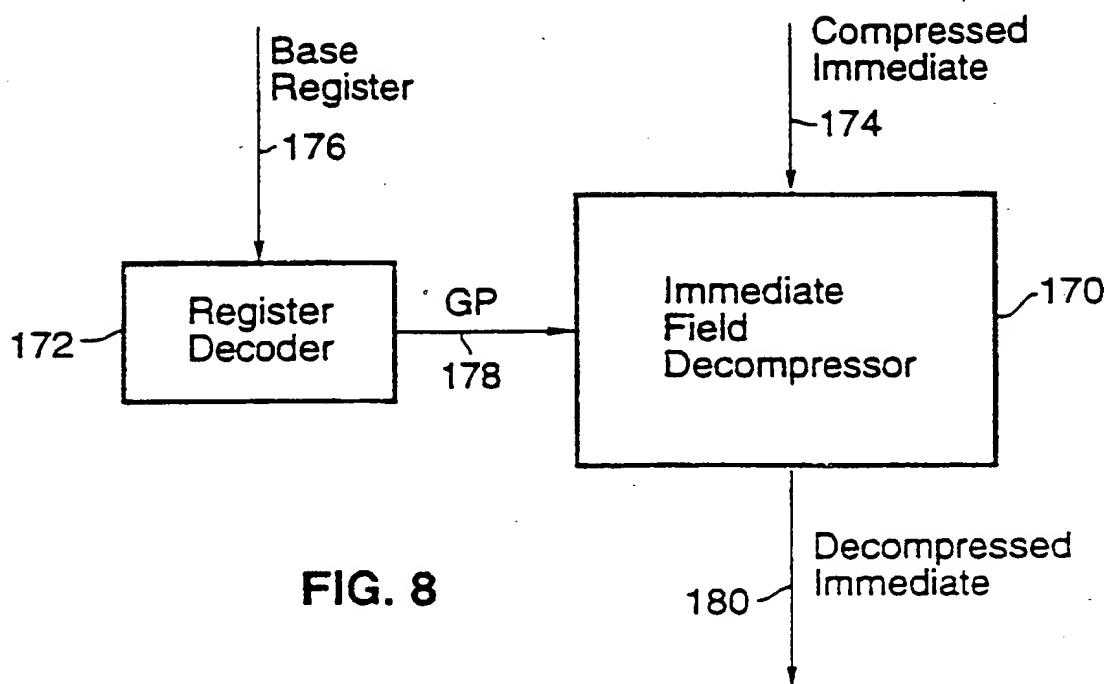


FIG. 7

From Instruction Cache 14B



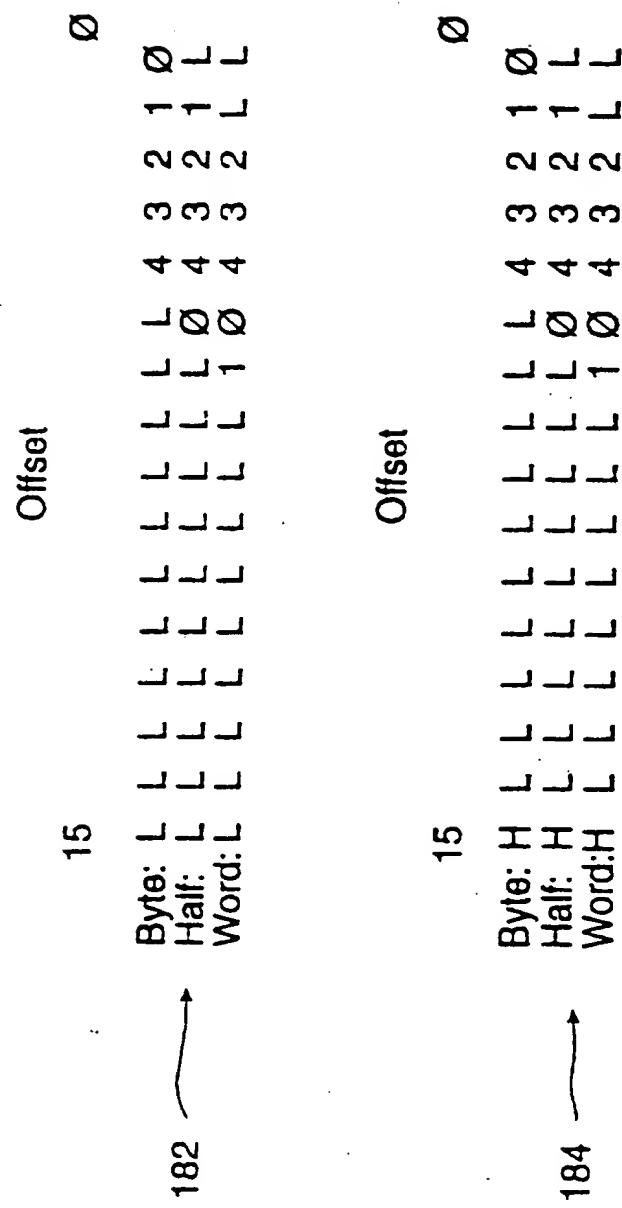


FIG. 9

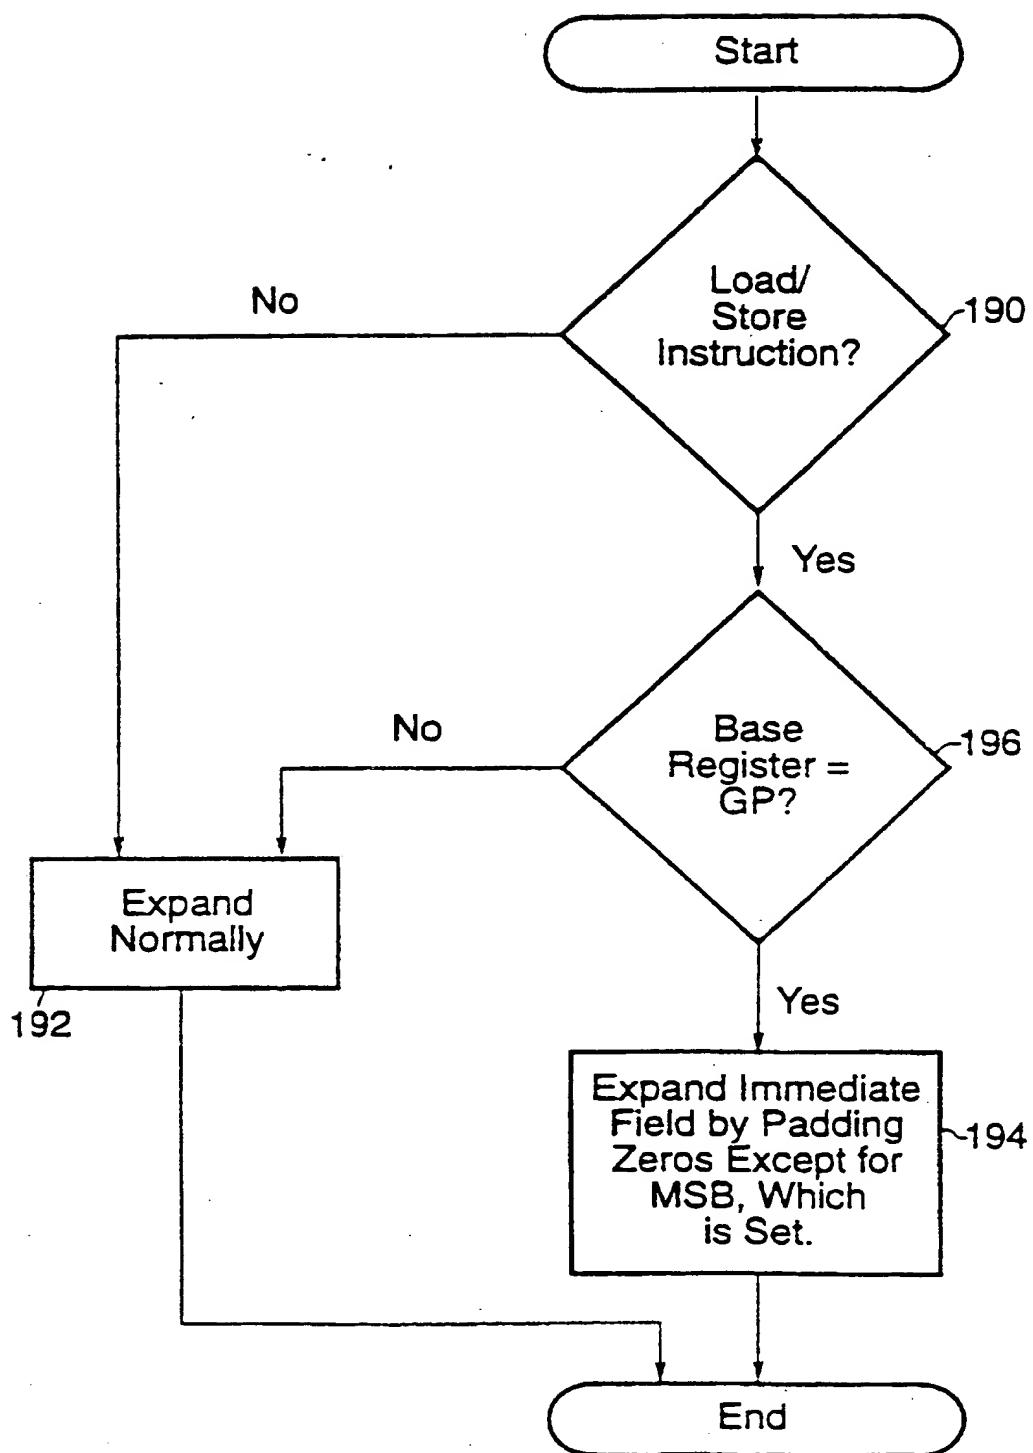


FIG. 10

From Instruction Cache 14B

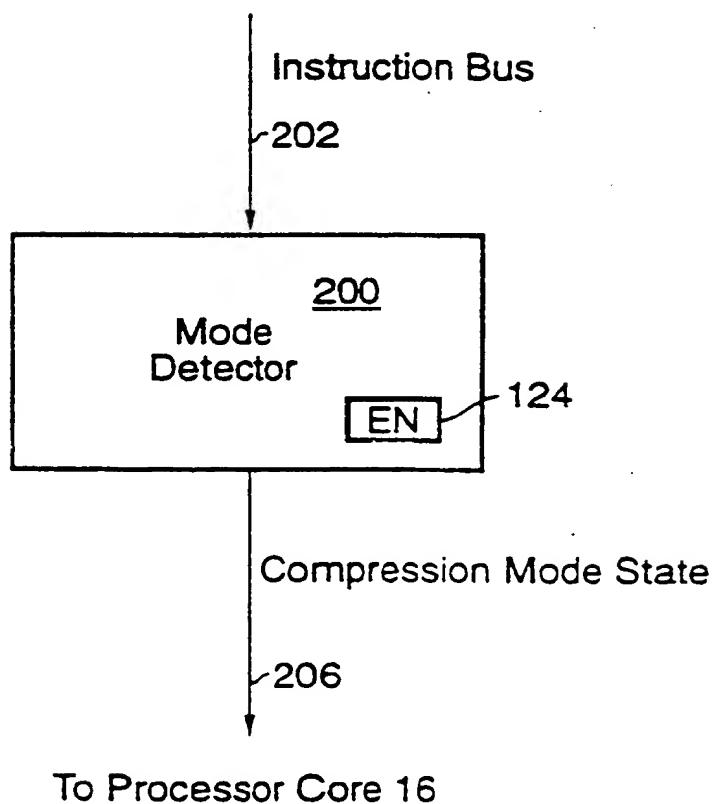


FIG. 11

From Instructor Cache 14B

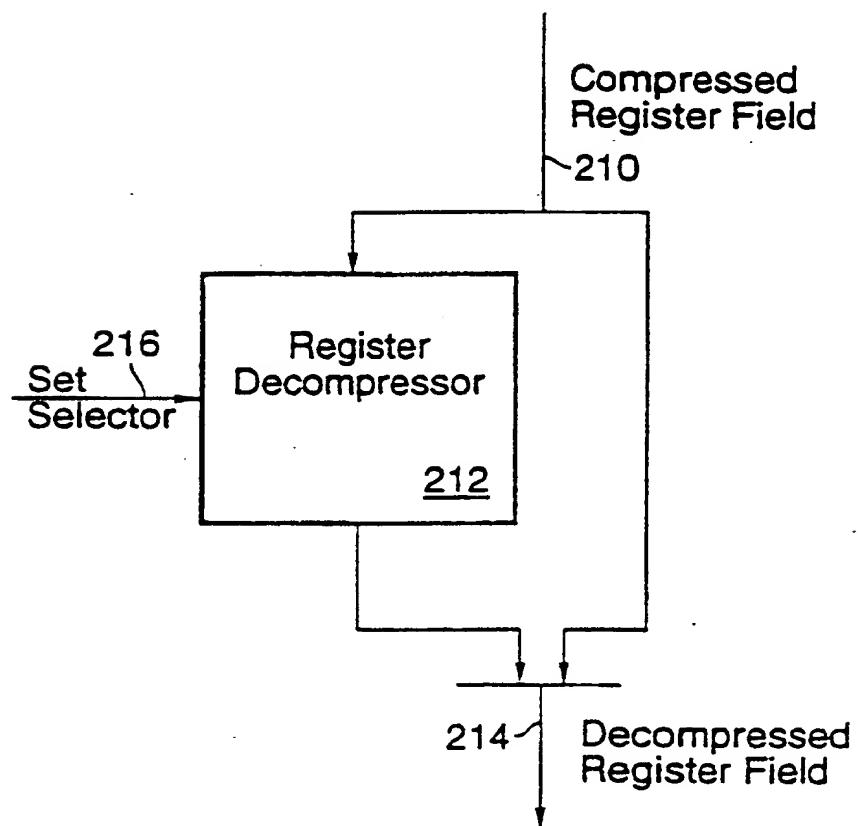


FIG. 12

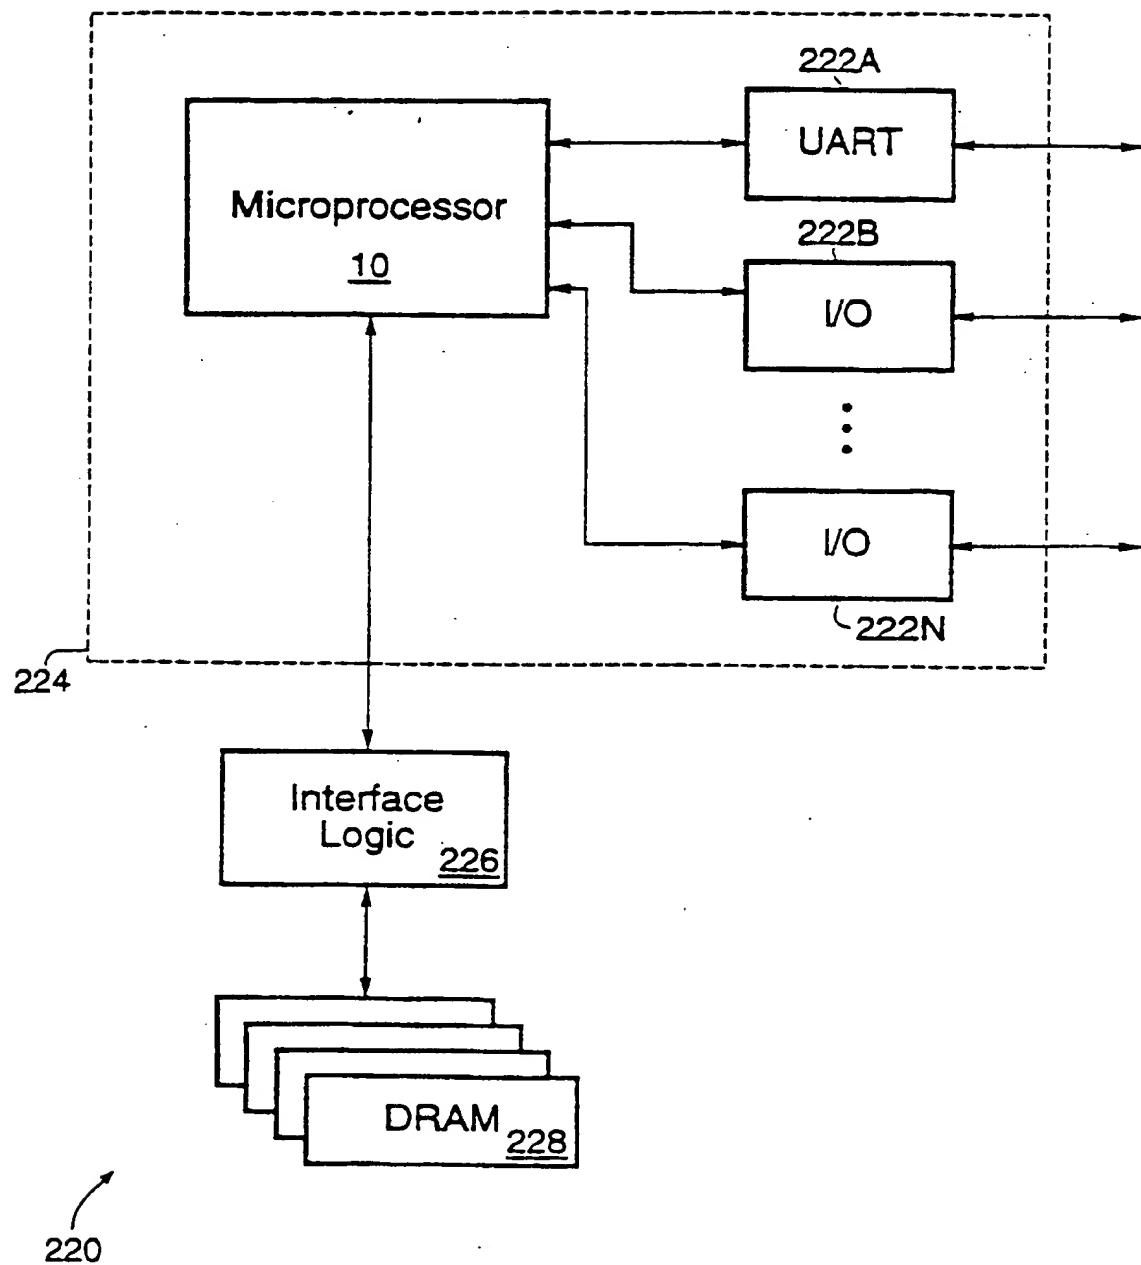


FIG. 13

TITLE: AN APPARATUS AND METHOD FOR DETECTING AND DECOMPRESSING INSTRUCTIONS FROM A VARIABLE-LENGTH COMPRESSED INSTRUCTION SET

5 BACKGROUND OF THE INVENTION

1. Field of the Invention

This invention relates to the field of microprocessors and, more particularly, to optimization of the
10 instruction set of a microprocessor.

2. Description of the Relevant Art

15 Microprocessor architectures may generally be classified as either complex instruction set computing (CISC) architectures or reduced instruction set computing (RISC) architectures. CISC architectures specify an instruction set comprising high level, relatively complex instructions. Often, microprocessors implementing CISC architectures decompose the complex instructions into multiple simpler operations which may be more readily implemented in hardware. Microcoded routines stored in an on-chip read-only memory (ROM) have been successfully employed for providing the decomposed operations corresponding to an instruction. More
20 recently, hardware decoders which separate the complex instructions into simpler operations have been adopted by certain CISC microprocessor designers. The x86 microprocessor architecture is an example of a CISC architecture.

Conversely, RISC architectures specify an instruction set comprising low level, relatively simple
25 instructions. Typically, each instruction within the instruction set is directly implemented in hardware. Complexities associated with the CISC approach are removed, allowing for more advanced implementations to be designed. Additionally, high frequency designs may be achieved more easily since the hardware employed to execute the instructions is simpler. An exemplary RISC architecture is the MIPS RISC architecture.

Although not necessarily a defining feature, variable-length instruction sets have often been associated with CISC architectures while fixed-length instruction sets have been associated with RISC architectures.
30 Variable-length instruction sets use dissimilar numbers of bits to encode the various instructions within the set as well as to specify addressing modes for the instructions, etc. Generally speaking, variable-length instruction sets attempt to pack instruction information as efficiently as possible into the byte or bytes representing each instruction. Conversely, fixed-length instruction sets employ the same number of bits for each instruction (the number of bits is typically a multiple of eight such that each instruction fully occupies a fixed number of bytes).
35 Typically, a small number of instruction formats comprising fixed fields of information are defined. Decoding each instruction is thereby simplified to routing bits corresponding to each fixed field to logic designed to decode that field.

Because each instruction in a fixed-length instruction set comprises a fixed number of bytes, locating
40 instructions is simplified as well. The location of numerous instructions subsequent to a particular instruction is implied by the location of the particular instruction (i.e. as fixed offsets from the location of the particular instruction). Conversely, locating a second variable-length instruction requires locating the end of the first variable-length instruction; locating a third variable-length instruction requires locating the end of the second variable-length instruction, etc. Still further, variable-length instructions lack the fixed field structure of fixed-length instructions. Decoding is further complicated by the lack of fixed fields.

Unfortunately, RISC architectures employing fixed-length instruction sets suffer from problems not generally applicable to CISC architectures employing variable-length instruction sets. Because each instruction is fixed length, certain of the simplest instructions may effectively waste memory by occupying bytes which do not convey information concerning the instruction. For example, fields which are specified as "don't care" fields for a particular instruction or instructions in many fixed-length instruction sets waste memory. In contrast, variable-length instruction sets pack the instruction information into a minimal number of bytes.

Still further, since RISC architectures do not include the more complex instructions employed by CISC architectures, the number of instructions employed in a program coded with RISC instructions may be larger than the number of instructions employed in the same program coded in with CISC instructions. Each of the more complex instructions coded in the CISC version of the program is replaced by multiple instructions in the RISC version of the program. Therefore, the CISC version of a program often occupies significantly less memory than the RISC version of the program. Correspondingly, more bandwidth between devices storing the program memory, and the microprocessor is needed for the RISC version of the program than for the CISC version of the program.

SUMMARY OF THE INVENTION

The problems outlined above are in large part solved by a microprocessor in accordance with the present invention. The microprocessor is configured to fetch a compressed instruction set which comprises a subset of a corresponding non-compressed instruction set. The non-compressed instruction set may be a RISC instruction set, such that the microprocessor may enjoy the high frequency operation and simpler execution resources typically associated with RISC architectures. Fetching the compressed instructions from memory and decompressing them within the microprocessor advantageously decreases the memory bandwidth required to achieve a given level of performance (e.g. instructions executed per second). Still further, the amount of memory occupied by the compressed instructions may be comparatively less than the corresponding non-compressed instructions may occupy.

The exemplary compressed instruction set described herein is a variable length instruction set. According to one embodiment, two distinct instruction lengths are included: 16-bit and 32-bit instructions. The 32-bit instructions are coded using an extend opcode, which indicates that the instruction being fetched is an extended (e.g. 32 bit) instruction. Instructions may be fetched as 16-bit quantities. When a 16-bit instruction having the extend opcode is fetched, the succeeding 16-bit instruction is concatenated with the instruction having the extend opcode to form a 32-bit extended instruction. Extended instructions have enhanced capabilities with respect to non-extended instructions, further enhancing the flexibility and power of the compressed instruction set. Routines which employ the capabilities included in the extended instructions may thereby be coded using compressed instructions.

The compressed instruction set further includes multiple sets of register mappings from the compressed register fields to the decompressed register fields. Each value coded in the compressed register fields decompresses to a different register within the microprocessor. In one embodiment, the compressed register fields comprise three bits each. Therefore, eight registers are accessible to a particular instruction. In order to offer access to additional registers for certain select instructions, the select instructions are assigned two opcode encodings. One of the opcode encodings indicates a first mapping of register fields, while the

second opcode encoding indicates a second mapping of register fields. Advantageously, the compressed register fields may include relatively few bits while select instructions for which access to additional registers is desired may be granted such access. Additionally, the register mappings are selected to minimize the logic employed to decompress register fields. In one embodiment, the compressed register field is directly copied 5 into a portion of the decompressed register field while the remaining portion of the decompressed register field is created using a small number of logic gates.

The microprocessor supports programs having routines coded in compressed instructions and other routines coded in non-compressed instructions. The subroutine call instruction within the compressed instruction set includes a compression mode which indicates whether or not the target routine is coded in 10 compressed instructions. The compression mode specified by the subroutine call instruction is captured by the microprocessor as the compression mode for the routine. In one embodiment, the compression mode is stored as one of the fetch address bits (stored in a program counter register within the microprocessor). Since the compression mode is part of the fetch address and the subroutine call instruction includes storing a return address for the subroutine, the compression mode of the calling routine is automatically stored upon execution 15 of a subroutine call instruction. When a subroutine return instruction is executed, the compression mode of the calling routine is thereby automatically restored.

An additional feature of one embodiment of the microprocessor is the decompression of the immediate field used for load/store instructions having the global pointer register as a base register. The immediate field is decompressed into a decompressed immediate field for which the most significant bit is set. A subrange of 20 addresses at the lower boundary of the global variable address space is thereby allocated for global variables of compressed instructions. Non-compressed instructions may store global variables in the remainder of the global variable address space. Advantageously, global variable allocation between the compressed and non-compressed routines of a particular program may be relatively simple since the subranges are separate.

Broadly speaking, the present invention contemplates an apparatus for executing instructions from a 25 variable-length compressed instruction set, comprising an instruction decompressor. The instruction decompressor is coupled to receive instructions which are members of the variable-length compressed instruction set, wherein the instruction decompressor is configured to examine an opcode field of a particular instruction. The instruction decompressor is configured to determine that the particular instruction is an extended instruction having a first fixed length if the opcode field is coded as an extend opcode. Additionally, the instruction decompressor is configured to determine that the particular instruction is a non-extended 30 instruction if the opcode field is coded as a second opcode different than the extend opcode.

The present invention further contemplates a method for expanding compressed instructions into 35 decompressed instructions. A compressed instruction is determined to be an extended instruction having a first fixed length if an opcode field of the compressed instruction is an extend opcode. If the opcode field of the compressed instruction is a second opcode different than the extend opcode, the compressed instruction is a non-extended instruction having a second fixed length. The compressed instruction is decompressed into a decompressed instruction. A number of bytes included in the compressed instruction is defined by the first fixed length if the compressed instruction is an extended instruction. Alternatively, the number of bytes is defined by the second fixed length if the compressed instruction is a non-extended instruction.

40 The present invention still further contemplates an apparatus for expanding compressed instructions into decompressed instructions, comprising a first determining means, a second determining means, and a decompressing means. The first determining means determines that a compressed instruction is an extended

instruction having a first fixed length if an opcode field of the compressed instruction is an extend opcode. The second determining means determines that the compressed instruction is a non-extended instruction having a second fixed length if the opcode field of the compressed instruction is a second opcode different than the extend opcode. The decompressing means decompresses the compressed instruction into a decompressed instruction. A number of bytes included in the compressed instruction is defined by the first fixed length if the compressed instruction is the extended instruction. Alternatively, the number of bytes is defined by the second fixed length if the compressed instruction is the non-extended instruction.

The present invention yet further contemplates a method for executing a program including a first routine and a second routine in a microprocessor. A subroutine call instruction is executed within the first routine, wherein the subroutine call instruction indicates that the second routine is to be executed via a target address of the subroutine call instruction. An indication within the subroutine call instruction is examined. If the indication is in a first state, the second routine is determined to be coded using compressed instructions. The second routine is determined to be coded using non-compressed instructions if the indication is in a second state different than the first state.

Furthermore, the present invention contemplates an apparatus for executing a program including a first routine and a second routine in a microprocessor, comprising an executing means and an examining means. The executing means executes a subroutine call instruction within the first routine. The subroutine call instruction indicates that the second routine is to be executed via a target address of the subroutine call instruction. The examining means examines an indication within the subroutine call instruction. The examining means determines that the second routine is coded using compressed instructions if the indication is in a first state. If the indication is in a second state, the examining means determines that the second routine is coded using non-compressed instructions.

The present invention still further contemplates an apparatus for fetching compressed and non-compressed instructions in a microprocessor, comprising a storage device and a mode detector. The storage device stores a compression enable indicator. Coupled to the storage device, the mode detector is configured to detect a compression mode of a target routine upon fetch of a subroutine call instruction specifying the target routine. The mode detector is configured to convey the compression mode to a processor core. The processor core is configured to fetch compressed instructions if the compression mode indicates compressed. Additionally, the processor core is configured to fetch non-compressed instructions if the compression mode indicates non-compressed.

The present invention yet further contemplates a microprocessor comprising an instruction decompressor and a processor core. The instruction decompressor is coupled to receive compressed instructions which are members of a variable-length compressed instruction set. The instruction decompressor is configured to decompress each received compressed instruction into a corresponding decompressed instruction. Coupled to receive decompressed instructions, the processor core is configured to execute the decompressed instructions.

The present invention additionally contemplates a method for executing instruction code. Compressed instructions are fetched, wherein the compressed instructions are members of a variable-length compressed instruction set. The compressed instructions are decompressed in an instruction decompressor, thereby forming corresponding decompressed instructions. The decompressed instructions are executed in a processor core.

The present invention still further contemplates an apparatus for executing instruction code.

comprising a fetching means, a decompressing means, and an executing means. The fetching means fetches compressed instructions which are members of a variable-length compressed instruction set. The decompressing means decompresses the compressed instructions, thereby forming corresponding decompressed instructions. The executing means executes the decompressed instructions.

5 Furthermore, the present invention contemplates an instruction decompressor configured to decompress compressed instructions. A first one of the compressed instructions is codable to access a first subset of registers defined for a corresponding non-compressed instruction set. Additionally, a second one of the compressed instructions is codable to access the first subset of registers and is further codable to access a second subset of registers.

10 The present invention further contemplates a method for decompressing compressed instructions. A particular compressed instruction having a first register field is decompressed using a first register mapping from compressed register indicators to decompressed register indicators if the particular compressed instruction is encoded using a first opcode. Alternatively, the particular compressed instruction having the first register field is decompressed using a second register mapping from compressed register indicators to decompressed register indicators if the particular compressed instruction is encoded using a second opcode.

15 The present invention still further contemplates an apparatus for decompressing compressed instructions comprising a decompressing means. The decompressing means is configured to decompress a particular compressed instruction having a first register field using a first register mapping from compressed register indicators to decompressed register indicators if the particular compressed instruction is encoded using a first opcode. Additionally, the decompressing means is configured to decompress the particular compressed instruction using a second register mapping from compressed register indicators to decompressed register indicators if the particular compressed instruction is encoded using a second opcode.

20 The present invention yet further contemplates an instruction decompressor configured to decompress a compressed register field of a compressed instruction into a decompressed register field of a decompressed instruction. A decompression of the compressed register field is dependent upon a first value coded into the compressed register field and a second value coded into an opcode field of the compressed instruction.

25 The present invention additionally contemplates a method for decompressing a compressed register field of a compressed instruction into a decompressed register field of a decompressed instruction. At least a portion of the compressed register field is directly copied into a portion of the decompressed register field. The remaining portion of the decompressed register field is produced by logically operating upon the compressed register field.

30 Moreover, the present invention contemplates an apparatus for decompressing a compressed register field of a compressed instruction into a decompressed register field of a decompressed instruction, comprising a first means and a second means. The first means is for directly copying at least a portion of the compressed register field into a portion of the decompressed register field. The first means is coupled to receive the compressed register field. Similarly coupled to receive the compressed register field, the second means is for logically operating upon the compressed register field to produce a remaining portion of the decompressed register field.

35 40 Furthermore, the present invention contemplates an instruction decompressor configured to decompress a compressed register field of a compressed instruction into a decompressed register field of a decompressed instruction. The instruction decompressor forms a first portion of the decompressed register field by copying at least a portion of the compressed register field thereto. Additionally, the instruction

decompressor includes a logic block which is configured to operate upon the compressed register field to produce a remaining portion of the decompressed register field.

BRIEF DESCRIPTION OF THE DRAWINGS

5 Other objects and advantages of the invention will become apparent upon reading the following detailed description and upon reference to the accompanying drawings in which:

Fig. 1 is a block diagram of one embodiment of a microprocessor.

10 Fig. 2 is a block diagram of a second embodiment of a microprocessor.

Fig. 3A is a first instruction format supported by one embodiment of the microprocessors shown in Figs. 1 and 2.

15 Fig. 3B is a second instruction format supported by one embodiment of the microprocessors shown in Figs. 1 and 2.

Fig. 3C is a third instruction format supported by one embodiment of the microprocessors shown in Figs. 1 and 2.

20 Fig. 3D is a fourth instruction format supported by one embodiment of the microprocessors shown in Figs. 1 and 2.

25 Fig. 4A is a fifth instruction format supported by one embodiment of the microprocessors shown in Figs. 1 and 2.

Fig. 4B is a sixth instruction format supported by one embodiment of the microprocessors shown in Figs. 1 and 2.

30 Fig. 4C is a seventh instruction format supported by one embodiment of the microprocessors shown in Figs. 1 and 2.

Fig. 4D is an eight instruction format supported by one embodiment of the microprocessors shown in Figs. 1 and 2.

35 Figs. 5A, 5B, 5C, 5D, and 5E are tables of exemplary instructions using the formats shown in Figs. 3A, 3B, 3C, and 3D.

40 Figs. 6A, 6B, 6C, 6D, 6E, and 6F are tables of exemplary instructions using the formats shown in Figs. 4A, 4B, 4C, and 4D.

Fig. 7 is a diagram depicting offsets from an arbitrary register and a global pointer register, according to one embodiment of the microprocessors shown in Figs. 1 and 2.

5 Fig. 8 is a block diagram of exemplary hardware for expanding an immediate field from a compressed instruction to a decompressed instruction.

10 Fig. 9 is a diagram depicting decompressed offsets in accordance with one embodiment of the microprocessors shown in Figs. 1 and 2.

15 Fig. 10 is a flow chart depicting operation of a decompressor for immediate fields according to one embodiment of the microprocessors shown in Figs. 1 and 2.

Fig. 11 is a block diagram of exemplary hardware for generating fetch addresses according to one embodiment of the microprocessors shown in Figs. 1 and 2.

15 Fig. 12 is a block diagram showing register decompression logic employed in one embodiment of the microprocessors shown in Figs. 1 and 2.

20 Fig. 13 is a block diagram of an exemplary computer system including the microprocessor for which embodiments are shown in Figs. 1 and 2.

25 While the invention is susceptible to various modifications and alternative forms, specific embodiments thereof are shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that the drawings and detailed description thereto are not intended to limit the invention to the particular form disclosed, but on the contrary, the intention is to cover all modifications, equivalents and alternatives falling within the spirit and scope of the present invention as defined by the appended claims.

DETAILED DESCRIPTION OF THE INVENTION

30 Turning now to Fig. 1, a block diagram of a first embodiment of a microprocessor 10A is shown. Microprocessor 10A includes an instruction decompressor 12A, an instruction cache 14A, and a processor core 16. Instruction decompressor 12A is coupled to receive instruction bytes from a main memory subsystem (not shown). Instruction decompressor 12A is further coupled to instruction cache 14A. Instruction cache 14A is coupled to processor core 16.

35 Generally speaking, microprocessor 10A is configured to fetch compressed instructions from the main memory subsystem. The compressed instructions are passed through instruction decompressor 12A, which expands the compressed instructions into decompressed instructions for storage within instruction cache 14A. Many of the compressed instructions occupy fewer memory storage locations than the corresponding decompressed instructions, advantageously reducing the amount of memory required to store a particular program. Additionally, since instructions are decompressed within microprocessor 10A, the bandwidth required to transport the compressed instructions from the main memory subsystem to microprocessor 10A is

reduced. Microprocessor 10A may be employed within a computer system having a relatively small main memory. Relatively large programs may be stored in the main memory due to the compression of instructions stored therein.

In one embodiment microprocessor 10A is configured to execute both compressed and non-compressed instructions on a routine-by-routine basis. In other words, a routine may be coded using either compressed instructions or non-compressed instructions. Advantageously, routines which may not be efficiently coded in the compressed instruction set may be coded using non-compressed instructions, while routines which are efficiently coded in the compressed instruction set are so coded. Microprocessor 10A may support a particular decompression of the immediate field for load/store instructions using the global pointer register as a base register, in order to support mixing of compressed and non-compressed instructions. The particular decompression is detailed further below. Additionally, a compression mode is detected by instruction decompressor 12A. The compression mode identifies the instruction set in which a routine is coded: compressed or non-compressed.

Instruction compression is achieved in microprocessor 10A by imposing certain limitations upon the available instruction encodings. By limiting the instruction encodings, instruction field sizes may be reduced (i.e. the number of bits within an instruction field may be decreased). For example, the number of available registers may be reduced to form the compressed instruction set. Because fewer registers are available, a smaller field may be used to encode the registers used as source and destination operands for the instruction. Instruction decompressor 12A expands the encoded register field into a decompressed register field. The decompressed register field is included in the decompressed instruction. The compressed instructions use the reduced instruction fields, thereby occupying less memory (i.e. fewer bits) than the original instruction encodings defined by the microprocessor architecture employed by processor core 16.

Instruction decompressor 12A is configured to accept compressed instructions and to decompress the instructions into the original instruction encodings. Each instruction field within a particular compressed instruction is expanded from the compressed field to a corresponding decompressed field within the corresponding decompressed instruction. The decompressed instruction is coded in the original instruction format supported by processor core 16.

Processor core 16 includes circuitry for fetching instructions from instruction cache 14A, decoding the instructions, and executing the instructions. The instructions supported by processor core 16 are specified by the microprocessor architecture employed therein. In one particular embodiment, processor core 16 employs the MIPS RISC architecture. However, it is understood that processor core 16 may employ any microprocessor architecture. Since instruction decompressor 12A decompresses instructions into the original instruction format, processor core 16 may comprise a previously designed processing core. In other words, the processing core may not require substantial modification to be included within microprocessor 10A.

The MIPS RISC architecture specifies an instruction set comprising 32 bit fixed-length instructions. A compressed instruction set is defined for microprocessor 10A which comprises variable-length instructions. Many of the compressed instructions comprise 16-bit instructions. Other compressed instructions comprised 32 bit instructions in conjunction with the extend instruction described below. Several 16-bit and 32-bit instruction formats are defined. It is understood that, although 16-bit and 32-bit compressed instructions are used in this embodiment, other embodiments may employ different instruction lengths. The compressed instructions encode a subset of the non-compressed instructions. Instruction encodings supported within the compressed instruction set comprise many of the most commonly coded instructions as well as the most often

used registers, such that many programs, or routines within the programs, may be coded using the compressed instructions.

In one embodiment, microprocessor 10A employs a compression mode. If the compression mode is active, then compressed instructions are being fetched and executed. Instruction decompressor 12A decompresses the instructions when they are transferred from main memory to instruction cache 14. Alternatively, the compressed mode may be inactive. When the compression mode is inactive, non-compressed instructions are being fetched and executed. Instruction decompressor 12A is bypassed when the compressed mode is inactive. In one particular embodiment, the compression mode is indicated by a bit within the fetch address (e.g. bit 0). The current fetch address may be stored in a PC register 18 within processor core 16. Bit 0 of PC register 18 indicates the compression mode (CM) of microprocessor 10A.

Instruction cache 14A is a high speed cache memory configured to store decompressed and non-compressed instructions. Although any cache organization may be employed by instruction cache 14A, a set associative or direct mapped configuration may be suitable for the embodiment shown in Fig. 1.

Turning next to Fig. 2, a second embodiment of a microprocessor 10B is shown. Microprocessor 10B includes an instruction cache 14B coupled to receive instruction bytes from the main memory subsystem, an instruction decompressor 12B, and processor core 16. Instruction cache 14B is coupled to instruction decompressor 12B, which is further coupled to processor core 16.

Microprocessor 10B is configured with instruction decompressor 12B between instruction cache 14B and processor core 16. Instruction cache 14B stores the compressed instructions transferred from the main memory subsystem. In this manner, instruction cache 14B may store a relatively larger number of instructions than a similarly sized instruction cache employed as instruction cache 14A in microprocessor 10A. Instruction decompressor 12B receives fetch addresses corresponding to instruction fetch requests from processor core 16, and accesses instruction cache 14B in response to the fetch request. The corresponding compressed instructions are decompressed into decompressed instructions by instruction decompressor 12B. The decompressed instructions are transmitted to processor core 16.

Similar to microprocessor 10A, microprocessor 10B includes a compression mode in one embodiment. Instruction decompressor 12B is bypassed when non-compressed instructions are being fetched and executed. For this embodiment, instruction cache 14B stores both compressed and non-compressed instructions. It is noted that instruction cache 14B typically stores instruction bytes in fixed-size storage locations referred to as cache lines. Therefore, a particular cache line may be storing compressed or non-compressed instructions. In either case, a plurality of instruction bytes are stored. Therefore, instruction caches 14A and 14B may be of similar construction. The compression mode at the time a cache line is accessed determines whether the instruction bytes are interpreted as compressed or non-compressed instructions.

An alternative configuration for microprocessor 10B is to include instruction decompressor 12B within the instruction decode logic of processor core 16. The compressed instructions may not actually be decompressed in such an embodiment. Instead, the compressed instructions may be decoded directly by the decode logic. The decoded instructions may be similar to the decoded instructions generated for the non-compressed instructions which correspond to the compressed instructions.

It is noted that microprocessors 10A and 10B are merely exemplary embodiments of a microprocessor 10 which operates upon compressed instructions. For the remainder of this discussion, microprocessor 10, instruction cache 14, and instruction decompressor 12 will be used to refer to the corresponding elements of

both Figs. 1 and 2 as well as other embodiments of the elements included in other implementations of microprocessor 10.

The terms decompression, compressed instruction, decompressed instruction, and non-compressed instruction are used in the above discussion and may further be used below. As used herein, the term "compressed instruction" refers to an instruction which is stored in a compressed form in memory. The compressed instruction is generally stored using fewer bits than the number of bits used to store the instruction when represented as defined in the microprocessor architecture employed by processor core 16. The term "decompressed instruction" refers to the result of expanding a compressed instruction into the original encoding as defined in the microprocessor architecture employed by processor core 16. The term "non-compressed instruction" refers to an instruction represented in the encoding defined by the microprocessor architecture employed by processor core 16. Non-compressed instructions are also stored in memory in the same format (i.e. non-compressed instructions were never compressed). Finally, the term "decompression" refers to the process of expanding a compressed instruction into the corresponding decompressed instruction.

It is noted that instruction decompressors 12A and 12B may be configured to simultaneously decompress multiple compressed instructions. Such embodiments of instruction decompressors 12 may be employed with embodiments of processor core 16 which execute multiple instructions per clock cycle.

Figs. 3A-3D and 4A-4D depict exemplary instruction formats for 16-bit and 32-bit compressed instructions, respectively, according to one specific embodiment of microprocessor 10 employing the MIPS RISC architecture. Other instruction formats may be employed by other embodiments. The instruction formats shown in Figs. 3A-3D each comprise 16 bits in this particular implementation. Conversely, the instruction formats shown in Figs. 4A-4D each comprise 32 bits in this particular implementation. The compressed instructions encoded using the instruction formats are decompressed into instruction formats as defined by the MIPS RISC architecture for each instruction.

Fig. 3A depicts a first instruction format 20. Instruction format 20 includes an opcode field 22, a first register field 24, a second register field 26, and a function field 28. Opcode field 22 is used to identify the instruction. Additionally, function field 28 is used in conjunction with certain particular encodings of opcode field 22 to identify the instruction. Effectively, function field 28 and opcode field 22 together form the opcode field for these instructions. When opcode field 22 employs certain other encodings than the particular encodings, function field 28 is used as an immediate field.

First register field 24 and second register field 26 identify destination and source registers for the instruction. The destination register is also typically used as a source register for the instruction. In this manner, two source operands and one destination operand are specified via first register field 24 and second register field 26. The notations "RT" and "RS" in first register field 24 and second register field 26 indicate the use of the fields in the instruction tables below. Either RT or RS may be a destination register, depending upon the encoding of the instruction.

In one embodiment, opcode field 22 comprises 5 bits, first register field 24 and second register field 26 comprise 3 bits each, and function field 28 comprises 5 bits. First register field 24 is divided into two subfields (labeled RT1 and RT0). RT1 comprises two bits in the present embodiment, while RT0 comprises one bit. RT1 is concatenated with RT0 to form first register field 24. Subfield RT1 and second register field 26 are used in certain instructions encoded via instruction format 20 to indicate one of the 32 registers defined by the MIPS RISC architecture.

Fig. 3B depicts a second instruction format 30. Instruction format 30 includes opcode field 22, first

register field 24, and second register field 26. Additionally, a third register field 32 and a function field 34 are shown. Third register field 32 is generally used to identify the destination register for instructions using instruction format 30. Therefore, first register field 24 and second register field 26 comprise source registers for instruction format 30. Function field 34 is used similar to function field 28. In the embodiment shown, third register field 32 comprises three bits and function field 34 comprises two bits.

A third instruction format 40 is shown in Fig. 3C. Instruction format 40 includes opcode field 22 and second register field 26, as well as an immediate field 42. Immediate field 42 is used to provide immediate data for the instruction specified by instruction format 40. Immediate data is an operand of the instruction, similar to the value stored in a register specified by first register field 24 or second register field 26. For example, an add instruction which uses immediate data adds the immediate data to the value stored in the destination register, and stores the resulting sum into that destination register. In one embodiment, immediate field 42 comprises eight bits. Immediate field 42 is divided into two subfields (IMM1 and IMM0) in the instruction format shown in Fig. 3C. The subfields allow second register field 26 to be placed in the same bit positions within instruction format 40 as it is placed in instruction formats 20 and 30. Advantageously, second register field 26 is always found in the same position of 16-bit instructions in which it is used. Therefore, subfield IMM1 comprises 2 bits and subfield IMM0 comprises 6 bits. IMM1 is concatenated with IMM0 to form the immediate value.

Fig. 3D depicts a fourth instruction format 50. Instruction format 50 includes opcode field 22 and an immediate field 52. Immediate field 52, similar to immediate field 42, is used as an operand of the instruction. However, immediate field 52 comprises 11 bits.

Fig. 4A depicts a fifth instruction format 60. Instruction format 60 includes opcode field 22, which is coded as the extend instruction. Instruction decompressor 12 recognizes the extend instruction opcode within opcode field 22 and treats the current instruction as a 32-bit instruction (i.e. the 16 bits included in the instruction containing the extend opcode and the 16 bits which would otherwise comprise the next instruction in program order are concatenated to form a 32 bit instruction). Therefore, the compressed instruction can be seen to be a variable-length instruction set comprising 16-bit instructions and 32-bit instructions. Instruction format 60 further includes a zero field 62 comprising six bits (coded to all binary zeros), an immediate field 64, and a BR field 66. Instruction format 60 is used to code an extended form of the BR instruction (an unconditional branch instruction), and hence BR field 66 is an opcode field indicating the BR instruction. In one embodiment, the BR opcode is hexadecimal 02.

The extended BR instruction has a larger immediate field than the non-extended BR instruction, and therefore may be coded with larger offsets than the non-extended BR instruction. When a branch to an instruction distant from the branch instruction is desired, the extended BR instruction may be used. Alternatively, branches to close instructions may use the non-extended BR instruction. Immediate field 64 comprises 16 bits which are used as an offset to be added to the address of the instruction following the BR instruction to create the target address of the branch instruction. The non-extended BR instruction, by contrast, includes an eleven bit offset (i.e. it is coded using instruction format 50).

Fig. 4B depicts an instruction format 70 which is an extended version of instruction format 40. Instruction format 70 includes opcode field 22 coded as the extend opcode, as well as an immediate field 72, a first register field 74, a second register field 76, and a second opcode field 78. First register field 74 and second register field 76 comprise five bits each in the embodiment shown. Therefore, any register defined by the MIPS RISC architecture may be accessed using instruction format 70. Second opcode field 78 defines the

instruction being executed, and comprises 5 bits (similar to opcode field 22). Finally, immediate field 72 comprises 12 bits divided into a one bit IMM2 subfield, a five bit IMM1 subfield, and a six bit IMM0 subfield. Immediate field 72 is formed by concatenating IMM2 with IMM1 and further with IMM0 in the embodiment shown.

5 An extended instruction format corresponding to instruction format 30 is shown in Fig. 4C as an instruction format 80. Instruction format 80 includes opcode field 22, first register field 74, second register field 76, and second opcode field 78, similar to instruction format 70. Additionally, instruction format 80 includes a third register field 82 and a function field 84. Third register field 82 is similar to third register field 32, except that third register field 82 comprises five bits. Therefore, any MIPS RISC architecture register may 10 be specified by third register field 82. Function field 84 is similar to function fields 28 and 34, except that function field 84 comprises six bits.

15 Second opcode field 78 is coded to a particular value to identify instruction format 80 from instruction format 70. When second opcode field 78 is coded to the particular value, instruction format 80 is assumed by instruction decompressor 12. Conversely, when second opcode field 78 is coded to a value other than the particular value, instruction format 70 is assumed by instruction decompressor 12. In one embodiment, the particular value comprises hexadecimal 00.

20 Instruction format 80 further includes a COP0 bit 86. COP0 bit 86, when set, indicates that certain coprocessor zero instructions (as defined in the MIPS RISC architecture) are being executed. The tables of 25 instructions below further define the instructions encoded by setting COP0 bit 86.

The instructions defined for instruction formats 20, 30, 40, and 50 are capable of performing many of the operations commonly performed in typical programs. However, routines may need to perform operations of which these instructions are incapable. While most of the instructions in the routine may be coded using instruction formats 20-50, several instructions may require additional encodings. For example, access to a register not included within the subset of available registers in formats 20-50 may be needed. Additional 30 instructions not included in the instructions encoded using formats 20-50 may be needed. For these and other reasons, the extend opcode and extended instruction formats 60-80 are defined.

35 Instruction decompressor 12 examines opcode field 22 in order to detect the extend opcode. The extend opcode is one of the opcodes defined to use instruction format 50 in the present embodiment although the bits included in immediate field 52 are assigned differing interpretations depending upon the extended instruction format coded for the particular extended instruction. The extended instruction formats include a second opcode field (e.g. fields 66 and 78) which identify the particular extended instruction.

40 Addition of the extend opcode and extended instruction formats allows for many instructions to be encoded using the narrower instruction formats 20-50, but still have the flexibility of the wider extended instruction formats when desired. Programs which occasionally make use of the functionality included in the extended instruction formats may still achieve a reduced memory footprint, since these programs may be encoded using compressed instructions and many of the compressed instructions may comprise 16-bit compressed instructions.

An embodiment of microprocessor 10 may handle the extended instructions by fetching 16-bit instruction portions and detecting the extend opcode. When the extend opcode is detected, a NOP may be transmitted to processor core 16 and the remaining 16-bit portion of the extended instruction may be fetched. The extended instruction is decompressed and provided as the next instruction after the NOP.

45 Additionally, instruction decompressor 12 handles cases wherein a portion of the extended instruction

is available while a second portion is unavailable. For example, two portions of the extended instruction may lie within two distinct cache lines within instruction cache 14. Therefore, one portion of the instruction may be fetched from instruction cache 14 while the other portion may not reside within instruction cache 14. The portion may then need to be stored within instruction decompressor 12 until the remaining portion is available.

5 Finally, Fig. 4D is an instruction format 90 used to explicitly expand the JAL instruction of the MIPS RISC instruction set. The JAL instruction is often used as a subroutine call instruction. Subroutines may be stored in memory at a great distance (address-wise) from the calling routine. Therefore, having the largest possible range of relative offsets (via an immediate field 92 comprising 26 bits) is important for the JAL instruction. Additionally, an exchange bit 94 is included in the instruction encoding. The exchange bit is used 10 to indicate the compressed/non-compressed nature of the instructions at the target address. If the bit is set, the target instructions are compressed instructions. If the bit is clear, the target instructions are non-compressed instructions. The value of exchange bit 94 is copied into bit 0 of the program counter within processor core 16. Bit 0 of the program counter may always be assumed to be zero, since the sixteen bit and thirty-two bit 15 instructions occupy at least two bytes each and instructions are stored at aligned addresses. Therefore, bit zero is a useful location for storing the compression mode of the current routine. Processor core 16 increments fetch addresses by 2 (instead of 4) when bit 0 is set, thereby fetching 16 bit compressed instructions through instruction decompressor 12.

Each instruction within the compressed instruction set employed by microprocessor 10 uses at least 20 one of the instruction formats shown in Figs. 3A-3D and Figs. 4A-4D. It is noted that opcode field 22 is included in each instruction format, and is located in the same place within each instruction format. The coding of opcode field 22 determines which instruction format is used to interpret the remainder of the instruction. A first portion of the opcode field encodings is assigned to instruction format 20; a second portion of the opcode field encodings is assigned to instruction format 30; etc.

25 As used herein, the term "instruction field" refers to one or more bits within an instruction which are grouped and assigned an interpretation as a group. For example, opcode field 22 comprises a group of bits which are interpreted as the opcode of the instruction. Additionally, first and second register fields 24 and 26 comprise register identifiers which identify a storage location within processor core 16 which store operands of the instruction. Additionally, the term immediate field refers to an instruction field in which immediate data is coded. Immediate data may provide an operand for an instruction. Alternatively, immediate data may be used 30 as an offset to be added to a register value, thereby producing an address. Still further, immediate data may be used as an offset for a branch instruction.

Figs. 5A-6F are tables listing an exemplary compressed instruction set for use by one particular implementation of microprocessor 10. The particular implementation employs the MIPS RISC architecture within processor core 16. Therefore, the instruction mnemonics listed in an instruction column 100 of the 35 tables correspond to instruction mnemonics defined in the MIPS RISC architecture (or defined for the instruction assembler, as described in "MIPS RISC Architecture" by Kane and Heinrich, Appendix D, Prentice Hall PTR, Upper Saddle River, New Jersey, 1992, incorporated herein by reference) with the following exceptions: CMPI, MOVEI, MOVE, NEG, NOT, and extend. These instructions translate to the following MIPS instructions (RS and RT refer to the 16-bit RS and RT):

| | | |
|------|-------|---------------|
| CMPI | XORI | S24, RS, imm8 |
| MOVI | ADDIU | RS, S0, simm8 |

| | | |
|--------------------------|-----|-------------|
| MOV | ADD | RS, \$0, RT |
| NEG | SUB | RS, \$0, RT |
| NOT | NOR | RS, \$0, RT |
| extend (described above) | | |

5

Additionally, the instruction tables use several symbols. In an operands column 102, the symbols rs, rt, xs, xt, and rd are used. Rs and xs refer to second register field 26 (or second register field 76), while rt and xt refer to first register field 24 (or first register field 74). Similarly, rd refers to third instruction field 32 (or third instruction field 82). As mentioned for one embodiment above, first register field 24, second register field 26, and third register field 32 comprise three bits each. Table 1 below lists the mapping of the field encodings (listed in binary) to registers in the MIPS RISC architecture for these symbols. Other mappings are also contemplated, as shown further below. Names assigned according to MIPS assembler convention are also listed in Table 1.

10

Table 1: Register Mappings

| <u>Field Encoding</u> | <u>RS, RT, RD</u> | <u>XS, XT</u> |
|-----------------------|-------------------|---------------|
| 000 | \$8 (t0) | \$24 (t8) |
| 001 | \$1 (at) | \$17 (s1) |
| 010 | \$2 (v0) | \$18 (s2) |
| 011 | \$3 (v1) | \$19 (s3) |
| 100 | \$4 (a0) | \$28 (gp) |
| 101 | \$5 (a1) | \$29 (sp) |
| 110 | \$6 (a2) | \$30 (s8) |
| 111 | \$7 (a3) | \$31 (ra) |

15

As shown in table 1, up to 16 registers are available for use in compressed instructions having registers fields 24, 26, or 32. Because each register field is three bits, only eight registers are available for a given opcode. Instructions which may access all sixteen registers are assigned two opcodes in the instruction tables below. Register selection is thereby a function of both a register field and opcode field 22. Advantageously, register fields may be encoded using fewer bits while still providing select instructions which may access a large group of registers.

20

Also listed in operands column 102 are symbols for the immediate fields 32, 42, 64, and 72. The symbol "imm" indicates an immediate field is included. If "imm" is preceded by an "s", the immediate field is signed and the decompression of the immediate field into the decompressed instruction is performed by sign extending the immediate field. If "imm" is not preceded by an "S", the immediate field is unsigned and immediate field decompression involves zero extending the immediate field. In one embodiment, immediate field decompression for load/store instructions comprises right rotation of the immediate bits by one bit for halfwords and two bits for words, followed by shifting of the immediate bits left by one bit for halfwords and two bits for words. Effectively, a seven bit immediate field is provided for words and a six bit immediate field for halfwords (in the 16-bit instruction formats). The MIPS RISC architecture defines that data addresses corresponding to load/store instructions are aligned for each instruction included in the exemplary compressed instruction set. Therefore, the least significant bit (for halfwords) and the second least significant bit (for words) may be set to zero. Bits in the compressed immediate field need not be used to specify these bits.

30

Finally, "imm" is post-fixed with a number indicating the number of bits included in the immediate field.

Opcode field 22 and function field 28 are decompressed as well. More particularly, opcode field 22 and function field 28 identify the instruction within the MIPS RISC architecture, in accordance with the tables shown in Figs. 5A-6F. The opcode and function fields of the decompressed instructions are coded in accordance with the MIPS RISC architecture definition.

Figs. 5A and 5B depict a table 110 and a table 112, respectively. Tables 110 and 112 list instructions from the exemplary compressed instruction set which use instruction format 20 shown in Fig. 3A. Instruction column 100 and operands column 102 are included, as well as an opcode column 106 and a function column 104. Opcode column 104 and function column 106 include hexadecimal numbers, and correspond to opcode field 22 and function field 28, respectively.

Table 110 includes several instructions which have an "imm5" coding in function column 104. The "imm5" coding appears for the load/store instructions within table 110, and indicates that function field 28 is used as an immediate field for these instructions. For other instructions, function field 28 is used in conjunction with opcode field 22 to identify a particular instruction within the compressed instruction set.

Additionally, opcode 1d is labeled as special in table 110. The special instructions have a specific interpretation of function field 28. In particular, if the most significant bit of the function field is clear, then the instruction is defined to be:

ADDIU rt, rs, simm4

20

wherein the "simm4" operand is formed from the remaining bits of function field 28. If the most significant bit of function field 28 is set, the instruction is defined to be:

ADDIU xl, xs, simm4

25

except for two special cases. If second register field 26 is coded to a zero, then the instruction is:

MOVEI xl, imm4

30 wherein again the imm4 operand is formed from the remainder of function field 28. Lastly, if second register field 26 is coded to 5 (hexadecimal), then the instruction is defined to be:

ADDIU sp, simm9

35 wherein the simm9 operand is formed from the remaining bits of function field 28 and first register field 24. The low order two bits of the simm9 operand are set to zero.

It is noted that the destination of the SLT and SLTU instructions shown in table 110 is the t8 register (register S24) according to one embodiment.

40 Table 112 shows an "imm3" and "imm6" operand for several instructions. The imm3 operand is coded into second register field 26, and the "imm6" operand is coded into both second register field 26 and first register field 24.

Additionally, table 112 includes the jump register (JR) instruction, having second register field 26 as

an operand. However, it is noted that in one embodiment subfield RT1 of first register field 24 is used in conjunction with second register field 26 to specify any of the MIPS RISC architecture registers for the JR instruction.

Turning now to Fig. 5C, a table 114 including instruction column 100, operands column 102, opcode column 106, and function column 104. Table 114 lists instructions from the exemplary instruction set which use instruction format 30 shown in Fig. 3B. Certain instructions within table 114 have hardcoded destination registers (i.e. the destination registers cannot be selected by the programmer, other than by using a different opcode). For these instructions, third register field 32 is combined with function field 34 to store the function field encoding shown in function column 104. Additionally, an instruction is shown which has an immediate operand in function column 104 and operands column 102. This instruction uses second register field 26 in conjunction with function field 34 to code the corresponding immediate field used by the instruction.

Figs 5D and 5E are tables 116 and 118 showing the instructions from the exemplary compressed instruction set which employ instruction formats 40 and 50, respectively. It is noted that the extend instruction is shown in table 118. However, the extend instruction actually indicates that the instruction is a 32-bit compressed instruction which uses one of instruction formats 60, 70, or 80.

Turning now to Figs. 6A and 6B, a table 120 and a table 122 are shown. Tables 120 and 122 depict those instructions from the exemplary compressed instruction set which are encoded using instruction format 70, shown in Fig. 4B.

Table 120 includes instruction column 100 and operands column 102, and further includes an opcode column 108. Opcode column 108 is similar to opcode column 106, except that the opcode encodings shown in opcode column 108 correspond to opcode field 78.

Table 122 includes an RT column 109 which corresponds to first register field 74. The coding of the RT field in the instructions shown in table 122 indicates which instruction is selected. The instructions shown in table 122 share a specific encoding in opcode field 78. In one embodiment the specific encoding is 00 (hexadecimal).

Figs. 6C, 6D, 6E, and 6F are tables 124, 126, 128, and 130 which depict instructions from the exemplary compressed instruction set which are encoded according to instruction format 80. Tables 124, 126, and 130 include a function column 107 which corresponds to encodings of function field 84. Table 128 includes an RS, RT column 105 which will be explained in more detail below.

Operands column 102 for table 124 includes immediate operands for certain instructions. The "imm5" operand is coded into second register field 76. The "imm15" operand is coded into a combination of first register field 74, second register field 76, and third register field 82.

The instructions listed in table 128 are identified via encodings of second register field 76, as shown in RS, RT column 105. Certain instructions are identified via second register field 76 in conjunction with first register field 74. Those instructions for which RS, RT column 105 includes an asterisk for the RT portion are identified via second register field 76, while those instructions for which RS, RT column 105 does not include an asterisk are identified by second register field 76 in conjunction with first register field 74. Instructions which are not identified via first register field 74 may use first register field 74 to encode an operand. The instructions listed in tables 128 and 130 are instructions for which COP0 bit 86 is set, while instructions listed in tables 124 and 126 are encoded with COP0 bit 86 clear.

Certain instructions in table 128 include an "imm6" operand. The "imm6" operand is coded into function field 84. Additionally, function field 84 is used to indicate the instructions shown in table 130 when

second register field 76 is coded to 1x (hexadecimal), wherein "x" indicates that the low order bits are don't cared.

Turning now to Fig. 7, a first addressing window 150 and a second addressing window 152 are shown according to one embodiment of microprocessor 10. At the center of addressing window 150 is the value of a base register (represented as Reg. on the left side of addressing window 150). The value of the base register identifies an address within the main memory subsystem. Addressing window 150 represents the range of addresses around the value of the base register which are accessible to a load/store instruction in the non-compressed instruction set according to one embodiment of the non-compressed instruction set. The non-compressed instruction set specifies that load/store instructions form the address of a memory operand via the sum of a value stored in a base register and a sixteen bit signed immediate field. In such an embodiment, the range of addresses has an upper boundary of 32767 greater than the base register and a lower boundary of 32768 less than the base register. Other embodiments may include larger or smaller ranges. As used herein, the term "base register" refers to a register which is specified by a load/store instruction as storing a base address, to which the signed immediate field is added to form the address of the memory operand operated upon by the instruction.

As shown in table 110, for example, load/store instructions within the 16-bit portion of the exemplary compressed instruction set include a five bit immediate field. This field is rotated right two bits and then shifted left two bits for word-sized memory operands, forming a seven bit immediate field (the largest of the immediate fields which may be formed using the five bits, according to one embodiment). The seven bit immediate field is then zero extended to form a positive offset from the base register in the corresponding decompressed instruction. A subrange 154 of addresses are therefore available for access by compressed instructions. Within addressing window 150, subrange 154 has an upper boundary of 127 greater than the base register and a lower boundary of the base register. However, subrange 154 may vary in size from embodiment to embodiment.

While subrange 154 may work well for many load/store instructions, a different subrange may be employed for use with the global pointer register. The global pointer register is a register assigned by software convention to locate an area of memory used for storing global variables. A global variable is a variable which is available for access from any routine within a program. In contrast, a local variable is typically accessible only to a particular routine or group of routines. In the MIPS instruction set, for example, register \$23 is often used as the global pointer register.

The area of memory around the global pointer register may therefore be viewed as a table of global variables. Each global variable is assigned an offset within the table. The offset corresponds to a particular immediate field value which may be added to the global pointer register in order to locate the global variable. For the embodiment shown in Fig. 7, for example, a 64 kilobyte table may be allocated for global variables as shown along the left side of addressing windows 150 and 152.

If compressed immediate fields are decompressed as described for addressing window 150, then the global variable table includes a section which is accessible to compressed instructions (corresponding to subrange 154) which is between two subranges 156 and 158 accessible to non-compressed instructions. As noted above, microprocessor 10 may support programs in which some routines are coded with non-compressed instructions while other routines are coded with compressed instructions. Allocating global variables in a particular program is complicated by the division of the non-compressed global variable subranges 156 and 158 of addressing window 150. Global variables may be allocated into subrange 158, for example, and then

global variable allocation must continue in subrange 156 (for non-compressed instructions). In other words, subrange 154 must be bypassed for global variables accessible to non-compressed instructions.

Microprocessor 10 may employ a decompression of the compressed immediate field for load/store instructions using the global pointer (GP) register which leads to addressing window 152. Addressing window 152 includes a subrange 160 accessible to compressed instructions and a subrange 162 accessible to non-compressed instructions. Advantageously, subrange 162 is a contiguous block of memory. Global variables for access by non-compressed instructions may be allocated into subrange 162, while global variables for access by compressed instructions may be allocated into subrange 160. Essentially, subrange 160 and subrange 162 form distinct tables of global variables for access by compressed and non-compressed instructions, respectively.

Addressing window 152 is achieved by decompressing the compressed immediate field as described-above, except that the most significant bit of the decompressed immediate field is set. If the compressed immediate field is coded with binary zeros, then the decompressed immediate field is 8000 (in hexadecimal). Since the decompressed immediate field is interpreted as a signed field for load/store instructions, the 8000 value is the most negative number available in the decompressed immediate field. Other encodings of the compressed immediate field are decompressed into negative numbers which form subrange 160. Subrange 160 forms the lower boundary of the range of addresses represented by addressing window 152 as shown in the embodiment of Fig. 7.

As used herein, the term memory operand refers to a value stored in a memory location within the main memory subsystem. Load instructions may be used to transfer the memory operand to a register within microprocessor 10. Conversely, store instructions may be used to transfer a value stored in a register into the memory operand storage location. A memory operand may be of various sizes (i.e. numbers of bytes). In one embodiment, three sizes are available: byte, halfword, and word. A halfword comprises two bytes, and a word comprises four bytes. Other memory operand sizes are contemplated for other embodiments.

Turning to Fig. 8, a block diagram of exemplary hardware within instruction decompressor 12 for decompressing the immediate field of a load/store instruction is shown. It is noted that multiple copies of the exemplary hardware shown in Fig. 8 may be employed to concurrently decompress multiple load/store instructions. The exemplary hardware shown in Fig. 8 is described in terms of microprocessor 10B. However, similar hardware may be employed within microprocessor 10A. The exemplary hardware includes a immediate field decompressor 170 and a register decoder 172.

When an instruction is conveyed to instruction decompressor 12B from instruction cache 14B, a portion of the instruction comprising the compressed immediate field for load/store instructions is conveyed to immediate field decompressor 170 upon a compressed immediate bus 174. For the exemplary instruction set described in Figs. 3A-6F, the compressed immediate field comprises function field 28 (shown in Fig. 3A). Additionally, the base register field for the compressed load/store instruction is conveyed upon a base register bus 176. For the exemplary instruction set shown in Figs. 3A-6F, the base register field comprises second register field 26.

Register decoder 172 decodes the register identified upon base register bus 176. If the base register is the global pointer register, register decoder 172 asserts a GP signal upon GP line 178 to immediate field decompressor 170. Otherwise, register decoder 172 deasserts the GP signal.

Immediate field decompressor 170 decompresses the compressed immediate field in one of two ways, dependent upon the GP signal. If the GP signal is deasserted, then immediate field decompressor 170 clears the most significant bit of the decompressed immediate field. Conversely, immediate field decompressor 170

sets the most significant bit of the immediate field if the GP signal is asserted. Therefore, a positive offset is created when a register other than the global pointer register is used as the base register. A negative offset is created when the global pointer register is used as the base register. Immediate field decompressor 170 conveys the decompressed immediate field upon a decompressed immediate bus 180.

Fig. 9 illustrates the decompressed immediate field generated for load/store instructions according to one embodiment of the exemplary compressed instruction set. The compressed immediate field of load/store instructions which do not employ the global pointer register as the base register are decompressed as indicated by reference number 182. The decompression for bytes, halfwords, and words are shown separately, with each bit position of the decompressed immediate field (or offset) represented by a numerical digit or an "L". Bits from the compressed immediate field are shown in the respective bit locations of the decompressed field via the numerical digits. The least significant bit of the compressed immediate field is represented by the digit 0, and the most significant bit of the compressed immediate field is represented by a 4. The letter "L" is used to indicate a bit position which is set to a binary zero.

Decompressed immediate fields corresponding to bytes, halfwords, and words for load/store instructions which use the global pointer register as a base register are indicated by reference number 184. Similar to the decompressed fields indicated by reference number 182, the decompressed fields indicated by reference number 184 depict numerals in bit positions which are filled with a bit from the compressed immediate field and the letter "L" is used to indicate a bit position which is set to a binary zero. Additionally, the most significant bit of each decompressed offset is set to a binary one (indicated by the letter "H").

Turning next to Fig. 10, a flow chart is shown depicting activities performed by instruction decompressor 12 in order to decompress instructions in accordance with the embodiment shown in Fig. 8. Although the steps shown in Fig. 10 are illustrated as serial in nature, it is understood that various steps may be performed in parallel.

Instruction decompressor 12 determines if a received instruction is a load/store instruction (decision block 190). If the instruction is not a load/store instruction, the instruction is expanded in accordance with a mapping between the compressed instructions (as illustrated in Figs. 3A-6F) and the corresponding decompressed instructions (step 192). If the instruction is a load/store instruction, then the base register specified by the instruction is examined (decision block 196). If the base register is the global pointer register, the immediate field is decompressed as indicated by reference number 184 in Fig. 9 (step 194). Alternatively, if the base register is not the global pointer register, the immediate field is decompressed as indicated by reference number 182 in Fig. 9 (step 192).

In addition to decompressing load/store offsets in a different manner for the global pointer register, microprocessor 10 also supports a compression mode for indicating which type of instructions are being executed by microprocessor 10 (i.e. compressed or non-compressed). Fig. 11 is a block diagram illustrating a portion of one embodiment of instruction decompressor 12. The illustrated portion determines the compression mode for each routine executed by microprocessor 10. The portion shown may be suitable for microprocessor 10B, and a similar portion may be employed by microprocessor 10A. Fig. 11 depicts a mode detector 200.

When an instruction is fetched by processor core 16, the instruction is received upon an instruction bus 202 by mode detector 200. Mode detector 200 detects when the jump and link (JAL) instruction is fetched, and further examines the exchange bit 94. If exchange bit 94 is set, the routine at the target address of the JAL instruction comprises compressed instructions. Therefore, the compression mode of the target routine is

compressed. Alternatively, exchange bit 94 may be clear. In this case, the compression mode of the target routine is uncompressed.

In addition to specifying the compression mode for the target routine, the JAL instruction causes the address of the instruction following the JAL instruction to be stored into register \$31 of the MIPS RISC architecture. This register may subsequently be used with the JR instruction to return from the target routine. Because the compression mode is stored as part of the address in this embodiment, the compression mode of the source routine is restored upon execution of the JR instruction. Advantageously, routines encoded in compressed instructions may be intermixed with routines encoded in non-compressed instructions. The new compression mode is conveyed to processor core 16 upon a compression mode line 206. It is noted that mode detector 200 may be included as a part of processor core 16 instead of instruction decompressor 12, in alternative embodiments.

The embodiment of mode detector 200 shown in Fig. 11 includes a storage 204 for a compression enable bit. If compression is enabled, the compression enable bit is set. When instructions are fetched in compressed mode and compression is enabled, instruction decompressor 12 decompresses the instructions. If the enable bit is clear, instruction compression is disabled for microprocessor 10. Instruction decompressor 12 is bypassed when instruction decompression is disabled. Furthermore, mode detector 200 indicates that the compression mode is non-compressed when instruction compression is disabled.

As used herein, a routine is an ordered set of instructions coded for execution by microprocessor 10. The routine may be coded in either compressed or non-compressed instructions, and is delimited by a subroutine call instruction and a return instruction. The delimiting subroutine call instruction is not included within the routine. Instead, the subroutine call instruction indicates the beginning of the routine via a target address included with the subroutine call instruction. The first instruction of the routine is stored at the target address. Additionally, the address of an instruction within the routine including the subroutine call instruction is saved so that a return instruction may be executed to return to the calling routine. In the exemplary compressed instruction set depicted in Figs. 3A-6F, the jal instruction may serve as a subroutine call instruction. Alternatively, the jalr instruction may serve as a subroutine call instruction.

A routine ends with a return instruction, which causes subsequent instruction execution to return to the address saved when the corresponding subroutine call instruction is executed. In other words, the target address of the return instruction is the saved address. For the exemplary compressed instruction set, the jr instruction may serve as a return instruction. Generally speaking, a target address is an address at which instruction fetching is to begin upon execution of the instruction corresponding to the target address.

Turning next to Fig. 12, a block diagram of one embodiment of register field decompression is shown. Other embodiments of register field decompression are contemplated. The compressed register field corresponding to an instruction is conveyed upon compressed register field bus 210. A register decompressor block 212 receives the compressed register field. Additionally, at least a portion of the compressed register field is incorporated into the decompressed register field which is then conveyed upon decompressed register field bus 214. The decompressed register field is thereby formed by concatenating at least a portion of the compressed register field to the value generated by register decompressor block 212.

In one embodiment, the entire compressed register field is concatenated into the decompressed register field. Additionally, the remaining portion of the decompressed register field depends upon which register set the instruction accesses (e.g. xs vs. rs and xt vs. rt). A set selector signal is received upon set selector bus 216 for each register, indicating whether the xs (xt) or the rs (rt) register set should be used. If the

set selector signal is asserted, then $xs(x_1)$ is selected. Otherwise, $rs(r_1)$ is selected. The set selector signal is asserted or deasserted based upon the opcode of the instruction being decompressed, in accordance with the exemplary compressed instruction set shown in Figs. 5A-6F. For example, the register mapping between compressed and decompressed registers shown in Table 1 may be employed. For such an example, register decompressor 212 may employ the following logic, wherein DR represents the decompressed register field, CR represents the compressed register field, and RH represents the corresponding set selector signal value:

$$DR[4:3] = \{RH, (RH \& CR[2] | !CR[2:0])\}$$

Several other register mappings are contemplated, examples of which are shown in tables 2-4 below, along with corresponding Verilog logic equations. It is noted that any register mapping may be employed by various embodiments of microprocessor 10.

Table 2: Second Exemplary Register Mappings

| <u>Field Encoding</u> | <u>RS, RT, RD</u> | <u>XS, XT</u> |
|-----------------------|-------------------|---------------|
| 000 | \$8 (t0) | \$24 (t8) |
| 001 | \$9 (t1) | \$25 (t9) |
| 010 | \$2 (v0) | \$18 (s2) |
| 011 | \$3 (v1) | \$19 (s3) |
| 100 | \$4 (a0) | \$28 (gp) |
| 101 | \$5 (a1) | \$29 (sp) |
| 110 | \$6 (a2) | \$30 (s8) |
| 111 | \$7 (a3) | \$31 (ra) |

$$DR[4:3] = \{RH, (RH \& CR[2] | !CR[2:1])\}$$

Table 3: Third Exemplary Register Mappings

| <u>Field Encoding</u> | <u>RS, RT, RD</u> | <u>XS, XT</u> |
|-----------------------|-------------------|---------------|
| 000 | \$16 (s0) | \$24 (t8) |
| 001 | \$1 (at) | \$9 (t1) |
| 010 | \$2 (v0) | \$10 (t2) |
| 011 | \$3 (v1) | \$11 (t3) |
| 100 | \$4 (a0) | \$28 (gp) |
| 101 | \$5 (a1) | \$29 (sp) |
| 110 | \$6 (a2) | \$30 (s8) |
| 111 | \$7 (a3) | \$31 (ra) |

$$DR[4:3] = \{(RH \& CR[2] | !CR[2:0]), RH\}$$

Table 4: Fourth Exemplary Register Mappings

| <u>Field Encoding</u> | <u>RS, RT, RD</u> | <u>XS, XT</u> |
|-----------------------|-------------------|---------------|
| 000 | \$16 (s0) | \$24 (t8) |
| 001 | \$17 (s1) | \$25 (t9) |
| 010 | \$2 (v0) | \$10 (t2) |
| 011 | \$3 (v1) | \$11 (t3) |

| | | |
|-----|----------|-----------|
| 100 | \$4 (a0) | \$28 (gp) |
| 101 | \$5 (a1) | \$29 (sp) |
| 110 | \$6 (a2) | \$30 (s8) |
| 111 | \$7 (a3) | \$31 (ra) |

$$DR[4:3] = ((RH \& CR[2] \mid\mid CR[2:1]), RH)$$

As indicated by the assembler assigned names shown in tables 1-4, various registers are assigned to various functions by software convention. For example, the MIPS assembler assigns the following meanings to registers:

5

Table 5: Software Convention for Register Names

| <u>Register</u> | <u>Software Name</u> | <u>Use</u> |
|------------------|----------------------|---|
| \$0 | none | Hardwired to zero |
| \$1 | \$at | used by assembler |
| \$2..\$3 | v0-v1 | Function results or static link |
| \$4..\$7 | a0-a3 | arguments for a subroutine |
| \$8..\$15 | t0-t9 | Temporary registers, not saved between subroutine calls |
| \$24..\$25 | | |
| \$16..\$23, \$30 | s0-s8 | Saved between subroutine calls |
| \$26..\$27 | k0-k1 | Reserved for operating system |
| \$28 | gp | Global Pointer |
| \$29 | sp | Stack Pointer |
| \$31 | ra | Return address |

It is desirable to provide access to both temporary and saved registers to routines coded in compressed instructions. Additionally, access to v0-v1, a0-a3, gp, sp, and ra are needed to operate with existing software.

10 The register mappings shown balance these qualities with the desire for register decompressor 212 to occupy a fairly small number of gates. Advantageously, a useful set of registers is selected from the MIPS register set while still maintaining a low gate count within register decompressor 212.

Turning now to Fig. 13, an exemplary computer system 220 including microprocessor 10 is shown. Many other computer systems employing microprocessor 10 are contemplated. Within computer system 220, microprocessor 10 is incorporated onto a semiconductor substrate 224 along with multiple I/O interfaces 222A-222N. The I/O interfaces interface to I/O devices external to substrate 224. An exemplary I/O interface 222A may be a universal asynchronous receiver/transmitter (UART).

20 Microprocessor 10 may be coupled to I/O interfaces 222 for communication therewith. Additionally, microprocessor 10 may be coupled to external interface logic 226, which further interfaces to one or more dynamic random access memory (DRAM) modules 228. DRAM modules 228 may store compressed and/or non-compressed instruction code, as well as data for use by the program represented by the compressed and/or non-compressed instruction code.

25 It is noted that the present discussion may refer to the assertion of various signals. As used herein, a signal is "asserted" if it conveys a value indicative of a particular condition. Conversely, a signal is "deasserted" if it conveys a value indicative of a lack of a particular condition. A signal may be defined to be asserted when it conveys a logical zero value or, conversely, when it conveys a logical one value.

Although a specific example of a compressed instruction set is shown and described herein, multiple variations, extensions, and modifications may be made to the exemplary compressed instruction set. These variations, extensions, and modifications are contemplated.

The following Verilog listing describes exemplary logic for instruction decompressor 12. Many different embodiments of the logic are contemplated, although the Verilog listing shown is one suitable example:

```

5      'timescale 1 ns / 1 ns
10     module tinyrisc_dp( xo, ci, dojal, ext, x );
15
15     'output[31:0]    xo;           // expanded instruction out
20     input[15:0]      ci;          // compressed instruction in
25     input            dojal;       // do jal
30     input            ext;         // extend
35     input[10:0]      x;          // extend bits
40
45     wire   ext_jal = ext | dojal;
50
55     /* ci[15:11] decodes */
60     wire   xsp    = ci[15:11] == 5'b00000;
65     wire   spany  = ~ext_jal & ci[15] & ~ci[14] & ~ci[11];
70     wire   splor2 = spany & ~ci[13];
75     wire   special = splor2 & ~ci[12];
80     wire   brjal  = ci[15:12] == 4'b0001;
85     wire   br    = brjal & ~ci[11];
90     wire   word   = ci[15] & (~ext | ~ci[13]) & ci[12] & ci[11];
95     wire   special3= ~ext_jal & ci[15:11] == 5'b11101;
100    wire   half   = ci[15] & (~ext | ~ci[13] | ~ci[12]) & ci[11] & ~special3;
105    wire   opi    = ~ci[15] & ci[14] & ~ci[12];
110    wire   opx    = ci[15:12] == 4'b0000 | ci[15] & ~ci[13] & ci[12] & ci[11];
115    wire   sll    = ~ext_jal & ci[15:11] == 5'b10110;
120    wire   rtx0   = ~ext_jal & (~ci[15] & ci[14] & ci[12] & ~(ci[13] & ci[11]));
125    wire   smx    = ~ci[15] & ~ci[14] & (ci[13] | ~ci[12]) | ci[15:11] == 5'b01001;
130    wire   special2= splor2 & ci[12];
135    wire   ximm   = ~(xsp | br);
140    wire   exti   = ext & ximm;
145    wire   x2z    = spany & ci[12];
150    wire   x4z    = spany | special3;
155    wire   xn    = ci[15:11] == 5'b01111;
160    wire   ill    = brjal | xn;
165    wire   rsza   = ~ext_jal & (ill | sll | ci[15:11] == 5'b01000);
170
175    /* & decodes */
180    wire   jr    = special & ci[4:1] == 4'b0100;
185    wire   ja1r  = jr & ci[0];
190    wire   negnot = special & ci[4] & ~ci[3] & ci[1] & ci[0];
195    wire   rseq0  = ci[8:6] == 3'b000;
200    wire   sp2x   = special2 & ~ci[0];
205    wire   s1t    = sp2x & ci[1];
210    wire   sp3x   = special3 & ci[4];
215    wire   sp3sp  = sp3x & ci[8:6] == 3'b101;
220
225    /* | decodes */
230    wire   i8    = ~ext_jal & ~ci[15] | sp3sp;
235    wire   i8s   = i8 & ~ill;
240    wire   rdrs  = sp2x & ~ci[1] | special & ci[4];
245    wire   rdrt  = special & ~(ci[4] | ci[3]) | ja1r;
250    wire   rrdt  = special2 & ci[0] | sll;

```

```

      wire shrs = ~dojal & ( brjal | ~ext & xn | ci[14:11] == 4'b0000 & ci[4:2] == 3'b00 & ( ext ? ~(ci[15] |
          ci[5] | x[10]) : ci[15] ) );
      wire rss = ext_jal | i11 | jr & ~ci[0];
      wire rsz = rsza | sp2x & ~ci[2] & ~ci[1] | negnot | sp3x & rseq0 | shrs ;
      wire rts = ~ext_jal & opi | sp3sp ;
      wire rot = ~( ext_jal | rts | jr | ~ci[15] & ~opi );
      wire snsP3 = ci[3] & special3 & ~(ci[4] & rseq0);
      wire sml = snsP3 & ~sp3sp;
      wire smm = ci[10] & ~ext_jal & srx | snsP3 ;
      wire smh = smm | ci[10] & ~ext_jal & br | x[10] & exti ;
      wire sml1 = exti & half ? ci[0] : smh ;
      wire sml2 = exti & word ? ci[1] : smh ;
      wire xs = sp2x & ci[4] | sp3x | opx ;
      wire xi = jalr | sp2x & ci[3] | sp3x;

15 assign xo[31] = ~dojal & ci[15] & ( ext | ~x4z );
assign xo[30] = ext & xsp & x[10];
assign xo[29] = ~dojal & ci[14] & ~( ext & ci[15] & ~ci[12] & ~ci[11] );
assign xo[28] = ~dojal & ( ci[13] & ~( spany | ~ext & word | special3 ) | ~ext & ~ci[15] & ~ci[14] | br );
assign xo[27] = dojal | ci[12] & ( ext | ~( spany | ci[15]&ci[14]&~ci[11] ) );
assign xo[26] = dojal | ci[11] ;

20 wire[4:0] rs = { rs5 ? ci[10:9] : {xs,(ci[8] & xs | rseq0)}, ci[8:6] };
assign xo[25:21] = rs & {5{~rsz}};

25 assign xo[20:16] = { xi, (ci[10] ? xi : ~(ci[9] | ci[5])), ci[10:9], ci[5] };
wire[4:0] rr = { xi, (ci[10] ? xi : ~(ci[9] | ci[5])), ci[10:9], ci[5] };

30 assign xo[20:16] = { rr & {5{rrr}} |
                        | rs & {5{rrs}} |
                        | x[9:5] & {5{ext_jal}} |
                        | {rxx0, rxx0, 3'b00}};

35 wire[4:0] rd = { 1'b0, !ci[4:2], ci[4:2] };
assign xo[15:11] = { rs & {5{rdrs}} |
                      | rr & {5{rdrt}} |
                      | rd & {5{rdrd}} |
                      | x[4:0] & {5{ext & ~ximm | dojal}} |
                      | { (2{slt | smh}), smh, sml2, sml1 } };

40 assign xo[10:6] = { rs & {((2{shrs}), {3{shrs|sll}})} |
                        | ci[15:11] & {5{dojal}} |
                        | x[4:0] & {5{exti}} |
                        | { ci[1]&sll, ci[0]&sll, 1'b0, ci[10]&i8s, ci[9]&i8s | ci[1]&~ext_jal&word } |
                        | { smm, smm, smm, sml, sml } };

45 assign xo[5] = ci[5] & (ext_jal | i8) | ci[4] & special |
                        | ci[0] & ~ext_jal & half | special2 | sml ;
assign xo[4] = ci[4] & ~x4z | ci[1] & sp3sp | sml ;
assign xo[3] = sp3sp ? ci[0] : ci[3] & ~x2z | sll ;
assign xo[2] = ci[2] & ~x2z ;
assign xo[1] = ci[1] & ~(~dojal & word | sll | sp3sp ) ;
assign xo[0] = ci[0] & ~(~dojal & half | sll | sp3sp ) | special2 & ci[2] ;

55 endmodule

```

In accordance with the above disclosure, a microprocessor has been described which executes instructions from both a compressed instruction set and a non-compressed instruction set. The microprocessor expands the compressed instructions into decompressed instructions or directly decodes the compressed instructions. Advantageously, routines coded using the compressed instruction set occupy a smaller amount of

memory than the corresponding routines coded in non-compressed instructions. Memory formerly occupied by such routines may be freed for use by other routines or data operated upon by such routines.

Numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.

5

CLAIMS:

1. An apparatus for executing instructions from a variable-length compressed instruction set comprising:
 - an instruction decompressor (12A, 12B) coupled to receive compressed instructions (20, 30, 40, 50, 60, 70, 80, 90) which are members of said variable-length compressed instruction set, wherein said instruction decompressor is configured to decompress said compressed instructions into decompressed instructions, and wherein said instruction decompressor is configured to examine a first opcode field (22) of a particular compressed instruction, and wherein said instruction decompressor is configured to determine that said particular compressed instruction is an extended instruction (60, 70, 80, 90) having a first fixed length if said first opcode field is coded as an extend opcode, and wherein said instruction decompressor is configured to determine that said particular compressed instruction is a non-extended instruction (20, 30, 40, 50) having a second fixed length if said first opcode field is coded as a second opcode different than said extend opcode, and wherein said extended instruction includes a second opcode field (66, 78), and wherein said instruction decompressor is configured to decompress said second opcode field of said extended instruction into an decompressed opcode field within a decompressed instruction, and wherein said instruction decompressor is configured to decompress said first opcode field of said non-extended instruction into said decompressed opcode field.
 2. The apparatus as recited in claim 1 wherein said first fixed length is a first number of bytes which is greater than a second number of bytes corresponding to said second fixed length.
 3. The apparatus as recited in claim 2 wherein said first number is greater than said second number by an integer factor.
 4. The apparatus as recited in claim 3 wherein said integer factor is two.
 5. The apparatus as recited in claim 1 wherein said instruction decompressor is configured to examine said second opcode field in order to determine how to decompress said extended instruction.
 6. The apparatus as recited in claim 1 further comprising a processor core (16) configured to execute said decompressed instructions.
 7. The apparatus as recited in claim 6 wherein each compressed instruction corresponds to one decompressed instruction.
 8. The apparatus as recited in claim 1, wherein a first one of said compressed instructions is codable to access

a first subset of registers defined for a corresponding non-compressed instruction set, and wherein a second one of said compressed instructions is codable to access said first subset of registers and is further codable to access a second subset of registers.

5 9. The apparatus as recited in claim 8 wherein said second one of said compressed instructions is assigned a first opcode encoding and a second opcode encoding.

10 10. The apparatus as recited in claim 9 wherein said first opcode encoding indicates that said second one of said compressed instructions is coded to access one of said first subset of registers.

10 11. The apparatus as recited in claim 10 wherein said instruction decompressor is configured to decompress said second one of said compressed instructions using a first mapping of compressed register codings to decompressed register codings, wherein said first mapping maps each compressed register coding to a decompressed register coding within said first subset.

15 12. The apparatus as recited in claim 11 wherein said second opcode encoding indicates that said second one of said compressed instructions is coded to access one of said second subset of registers.

20 13. The apparatus as recited in claim 12 wherein said instruction decompressor is configured to decompress said second one of said compressed instructions using a second mapping of compressed register codings to decompressed register codings, wherein said second mapping maps each of said compressed register codings to a decompressed register coding within said second subset of registers.

25 14. The apparatus as recited in claim 9 wherein said second one of said compressed instructions includes a first register field (26) and a second register field (32).

15 15. The apparatus as recited in claim 14 wherein said instruction decompressor decompresses said first register field according to said first subset of registers if said first opcode encoding is used.

30 16. The apparatus as recited in claim 15 wherein said instruction decompressor decompresses said first register field according to said second subset of registers if said second opcode encoding is used.

35 17. The apparatus as recited in claim 16 wherein said second one of said compressed instructions is assigned a third opcode encoding, wherein said decompressor decompresses said second register field using said second subset of registers if said third opcode encoding is used.

18. The apparatus as recited in claim 17 wherein said second one of said compressed instructions is assigned a fourth opcode encoding, wherein said decompressor decompresses said first register field and said second

register field using said second subset of registers if said fourth opcode encoding is used.

19. The apparatus as recited in claim 9 wherein said first opcode encoding and said second opcode encoding differ in bits included in a function field (28, 34) of said second one of said compressed instructions.

5 20. The apparatus as recited in claim 1, wherein said instruction decompressor is configured to decompress a compressed register field (26, 32) of one of said compressed instructions into a decompressed register field of a corresponding decompressed instruction responsive to a first value coded into said compressed register field and a second value coded into said first opcode field of said compressed instruction.

10 21. The apparatus as recited in claim 20 wherein said instruction decompressor is configured to select one of multiple mappings from compressed register field encodings to decompressed register field encodings in response to said second value.

15 22. The apparatus as recited in claim 21 wherein said instruction decompressor is configured to select one of said decompressed register encodings from said one of said multiple mappings in response to said first value.

20 23. The apparatus as recited in claim 22 wherein a particular instruction is assigned a first opcode encoding and a second opcode encoding, and wherein said instruction decompressor, upon receipt of said particular instruction having said first opcode encoding selects said one of said multiple mappings for decompressing said compressed register field.

25 24. The apparatus as recited in claim 23 wherein said instruction decompressor, upon receipt of said particular instruction having said second opcode encoding, is configured to select another one of said multiple mappings for decompressing said compressed register field.

30 25. The apparatus as recited in claim 1 wherein said instruction decompressor is configured to decompress a compressed register field (26, 32) of a compressed instruction into a decompressed register field of a decompressed instruction, wherein said instruction decompressor forms a first portion of said decompressed register field by copying at least a portion of said compressed register field thereto, and wherein said instruction decompressor includes a logic block (212) which is configured to operate upon said compressed register field to produce a remaining portion of said decompressed register field.

35 26. The apparatus as recited in claim 25 wherein said portion of said compressed register field comprises an entirety of said compressed register field.

27. The apparatus as recited in claim 25 wherein said logic block is configured to select a register mapping from compressed register field encodings to decompressed register field encodings in response to a signal

received from said instruction decompressor.

28. The apparatus as recited in claim 27 wherein said instruction decompressor is configured to assert said signal upon detection of a first opcode encoding within said first opcode field, said first opcode encoding assigned to a particular instruction, and wherein said instruction decompressor is configured to deassert said signal upon detection of a second opcode encoding within said first opcode field, said second opcode encoding assigned to said particular instruction.

29. A method for expanding compressed instructions into decompressed instructions, comprising:

10 determining that a compressed instruction (20, 30, 40, 50, 60, 70, 80, 90) is an extended instruction (60, 70, 80) having a first fixed length if an opcode field (22) of said compressed instruction is an extend opcode, wherein a number of bytes included in said compressed instruction is defined by said first fixed length if said compressed instruction is said extended instruction, and wherein said extended instruction includes a second opcode field (66, 78);

15 determining that said compressed instruction is a non-extended instruction (20, 30, 40, 50) having a second fixed length if said opcode field of said compressed instruction is a second opcode different than said extend opcode, wherein said number of bytes is defined by said second fixed length if said compressed instruction is said non-extended instruction; and

20 decompressing said compressed instruction into a decompressed instruction, said decompressed instruction including a decompressed opcode field, wherein said decompressing includes decompressing said second opcode field into said decompressed opcode field responsive to said determining that said compressed instruction is said extended instruction, and wherein said decompressing includes decompressing said opcode field into said decompressed opcode field responsive to said determining that said compressed instruction is said non-extended instruction.

25 30. The method as recited in claim 29 further comprising fetching compressed instructions in quantities defined by said second fixed length.

30 31. The method as recited in claim 30 further comprising forming said extended instruction, when detected, from a first of said quantities in which said extend opcode is detected and a second of said quantities which is immediately subsequent to said first of said quantities.

35 32. The method as recited in claim 29 further comprises identifying a particular extended instruction using said second opcode field.

33. The method as recited in claim 32 wherein said decompressing comprises identifying an extended instruction format (60, 70, 80) corresponding to said particular extended instruction.

5 34. The method as recited in claim 33 wherein said decompressing further comprises interpreting bits within said particular extended instruction according to a plurality of instruction fields within said extended instruction format.

10 35. The method as recited in claim 29 wherein said decompressing comprises identifying a non-extended instruction format (20, 30, 40, 50) corresponding to said non-extended instruction from an encoding of said opcode field.

15 36. The method as recited in claim 35 wherein said decompressing further comprises interpreting bits from said non-extended instruction according to a second plurality of instruction fields within said non-extended instruction format.

37. The method as recited in claim 29, further comprising:

20 decompressing a particular compressed instruction having a first register field (26, 32) using a first register mapping from compressed register indicators to decompressed register indicators for decompressing said first register field if said particular compressed instruction is encoded using a first opcode; and

25 decompressing said particular compressed instruction having said first register field using a second register mapping from compressed register indicators to decompressed register indicators for decompressing said first register field if said particular compressed instruction is encoded using a second opcode.

30 38. The method as recited in claim 37 wherein said particular compressed instruction further includes a second register field (26, 32).

39. The method as recited in claim 38 further comprising decompressing said second register field using said second register mapping if said particular compressed instruction is encoded using a third opcode.

35 40. The method as recited in claim 39 further comprising decompressing said second register field and said first register field using said second register mapping if said particular compressed instruction is encoded using a fourth opcode.

41. The method as recited in claim 37 further comprising decompressing a second particular compressed instruction using said first mapping.

42. The method as recited in claim 29, further comprising:

directly copying at least a portion of a compressed register field (26, 32) within one of said compressed instructions into a portion of a decompressed register field of a corresponding decompressed instruction; and

logically operating upon said compressed register field to produce a remaining portion of said decompressed register field.

43. The method as recited in claim 42 wherein said portion of said compressed register field comprises an entirety of said compressed register field.

44. The method as recited in claim 43 wherein said portion of said decompressed register field receiving said entirety of said compressed register field comprises a plurality of least significant bits of said decompressed register field.

45. The method as recited in claim 42 further comprising logically operating upon an opcode field (22) of said compressed instruction to produce said remaining portion of said decompressed register field.

46. The method as recited in claim 45 wherein said logically operating upon said opcode field comprises selecting a first register mapping in response to a first opcode encoding in said opcode field.

47. The method as recited in claim 46 wherein said logically operating upon said opcode field further comprises selecting a second register mapping in response to a second opcode encoding in said opcode field.

48. The method as recited in claim 47 wherein said first opcode encoding and said second opcode encoding are assigned to a particular instruction.

49. An apparatus substantially as hereinbefore described as an embodiment and as shown in the corresponding accompanying drawings.

50. A method substantially as hereinbefore described as an embodiment with reference to the corresponding accompanying drawings.

5



INVESTOR IN PEOPLE

Application No: GB 0018379.8
Claims searched: 1-48

Examiner: Ben Micklewright
Date of search: 15 August 2000

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.R): G4A (AJR APX)

Int Cl (Ed.7): G06F (9/30 9/318)

Other:

Documents considered to be relevant:

| Category | Identity of document and relevant passage | Relevant to claims |
|----------|--|--------------------|
| A | GB2289353 A (ADVANCED RISK MACH LTD) | - |
| A | GB2284492 A (SMITH GRAEME ROY) | - |
| A | US4437149 A (POMERENE JAMES H ET AL) | - |
| A | JP070121352 (CANON) | - |
| A | 14 th Annual International Symposium on Computer Architecture. Conference proceedings, pages 309-319, 1987, IEEE, D R Ditzel et al, "The hardware architecture of the CRISP microprocessor" | - |

| | |
|---|--|
| X Document indicating lack of novelty or inventive step | A Document indicating technological background and/or state of the art |
| Y Document indicating lack of inventive step if combined with one or more other documents of same category. | P Document published on or after the declared priority date but before the filing date of this invention. |
| & Member of the same patent family | E Patent document published on or after, but with priority date earlier than, the filing date of this application. |